Frequency-resolved Lags in UV/Optical Continuum Reverberation Mapping

Edward M. Cackett[1] , Abderahmen Zoghbi[2] , and Otho Ulrich[1,3]
[1] Wayne State University, Department of Physics & Astronomy, 666 W Hancock St, Detroit, MI 48201, USA; ecackett@wayne.edu
[2] Department of Astronomy, University of Michigan, 1085 South University Avenue, Ann Arbor, MI 48109, USA
[3] Western Michigan University, Department of Physics, 1120 Everett Tower, Kalamazoo, MI 49008, USA
Received 2021 September 5; revised 2021 October 29; accepted 2021 November 10; published 2022 January 21

Abstract

In recent years, continuum-reverberation mapping involving high-cadence UV/optical monitoring campaigns of nearby active galactic nuclei has been used to infer the size of their accretion disks. One of the main results from these campaigns has been that in many cases the accretion disks appear too large, by a factor of 2–3, compared to standard models. Part of this may be due to diffuse continuum emission from the broad-line region (BLR), which is indicated by excess lags around the Balmer jump. Standard cross-correlation lag-analysis techniques are usually used to just recover the peak or centroid lag and cannot easily distinguish between reprocessing from the disk and BLR. However, frequency-resolved lag analysis, where the lag is determined at each Fourier frequency, has the potential to separate out reprocessing on different size scales. Here we present simulations to demonstrate the potential of this method and then apply a maximum-likelihood approach to determine frequency-resolved lags in NGC 5548. We find that the lags in NGC 5548 generally decrease smoothly with increasing frequency, and are not easily described by accretion-disk reprocessing alone. The standard cross-correlation lags are consistent with lags at frequencies lower than 0.1 day^{-1}, indicating they are dominated from reprocessing at size scales greater than $\sim$10 light days. A combination of a more distant reprocessor, consistent with the BLR, along with a standard-sized accretion disk is more consistent with the observed lags than a larger disk alone.

Unified Astronomy Thesaurus concepts: Active galactic nuclei (16); Supermassive black holes (1663); Reverberation mapping (2019); Seyfert galaxies (1447); Accretion (14)

1. Introduction

The accretion disk (Shakura & Sunyaev 1973), X-ray corona (Haardt & Maraschi 1991), and other regions surrounding the central supermassive black hole in active galactic nuclei (AGNs) are beyond current spatial-resolution limits for the vast majority of objects. To overcome this, the reverberation-mapping technique uses time lags between light curves observed at different wavelengths to infer the geometry and kinematics of these inner regions (see Cackett et al. 2021, for a recent review). The time lags are assumed to be due to light travel time, with the path length to the observer different for different regions. While reverberation mapping was first used to determine the size scale of the Hβ-emitting broad-line region (BLR), more recently there has been significant effort to measure the size and temperature profile of accretion disks.

If a central ionizing source (often assumed to be the X-ray corona) irradiates the accretion disk, then variations in the irradiating source can drive variations in the heating of the disk where it is thermally reprocessed and reemitted. The hotter, inner part of the disk will see the variations before the cooler outer part of the disk, leading to correlated continuum light curves with the UV leading the optical. A standard Shakura & Sunyaev (1973) accretion disk with temperature profile $T \propto R^{-3/4}$ will give rise to wavelength-dependent lags, τ, following $\tau \propto \lambda^{4/3}$ (see Cackett et al. 2007, and references therein, for a more detailed discussion).

UV/optical continuum light curves have long been known to be well correlated, and to have short interband lags of a few days, at most (e.g., Stripe et al. 1994; Wanders et al. 1997;

Collier et al. 1998; Sergeev et al. 2005). Recent high-cadence (i.e., better than once per day) observations combining Swift and ground-based observations have significantly improved the continuum-lag measurements (Edelson et al. 2015, 2019; Fausnaugh et al. 2016; McHardy et al. 2018; Cackett et al. 2018, 2020; Hernández Santisteban et al. 2020; Kara et al. 2021). These campaigns have found common results: that the lags approximately follow $\tau \propto \lambda^{4/3}$; that the normalization of this relation is larger than expected for a standard disk; that lags in the u and U bands are generally systematically longer than expected based on an extrapolation of the other UV/optical lags; and that there is not a consistent relation between the X-rays and the UV/optical.

That the lags in the u/U band are offset by approximately a factor of 2 (Edelson et al. 2019), was further highlighted by Hubble Space Telescope (HST) spectroscopic monitoring of NGC 4593, which resolved the continuum lags in this region of the spectrum, showing a clear discontinuity at the Balmer jump (Cackett et al. 2018). These results can be explained if there is significant additional continuum emission not from the disk but from the BLR (Korista & Goad 2001, 2019; Lawther et al. 2018; Netzer 2020, 2022). This diffuse continuum from the BLR should emit across the full UV/optical bands, and can affect the lags in all wavelengths. The lag from this component increases with wavelength, except for discontinuities at the Balmer and Paschen jumps, and it is therefore hard to cleanly separate out its contribution (Korista & Goad 2001).

The spatial separation of the accretion disk and BLR-continuum-emitting regions should lead to the signals being distinguishable on different timescales, with disk reverberation taking place on timescales of a few days and BLR continuum reverberation taking place on timescales of weeks or longer. In principle, filtering the light curves on these timescales can be

used to pick out the different signals. For instance, techniques such as subtracting off a moving box-car average have been used in some cases to remove long-term variability, essentially acting as a high-pass filter (e.g., McHardy et al. 2014, 2018; Pahari et al. 2020; Vincentelli et al. 2021). Other approaches are to implement more sophisticated time-lag analysis. For instance, one can use maximum entropy techniques to try and recover the response function through fitting the light curves (Horne 1994). This has been used to show that in NGC 4593 the response function has both a prompt response and a tail to longer lags, which could be interpreted as signals from both the disk and the BLR (McHardy et al. 2018). Alternatively, Chelouche et al. (2019) use a bivariate reverberation model to suggest that the lags in Mrk 279 are dominated by the BLR.

Frequency-resolved lag analysis is yet another approach. This is widely used to perform X-ray reverberation (see, for a review, Uttley et al. 2014), where it has been able to successfully separate out lags from different processes taking place on different timescales (e.g., Kara et al. 2013). However, in X-ray analysis, typically the observation lengths are significantly longer than the timescale of interest, providing a continuous time series that can be relatively simply analyzed with a fast Fourier transform to get lags at different frequencies. The case in UV/optical reverberation is, of course, different with time series built up over many months but sampled irregularly, and often with gaps due to poor weather at the ground-based observatories. However, several techniques have been developed in order to deal with gaps in light curves and still perform a frequency-resolved analysis. Not all X-ray satellites provide long continuous exposures, thus Zoghbi et al. (2013) followed the method described by Miller et al. (2010) to develop a maximum-likelihood approach to determining power spectra and lags from unevenly sampled light curves. Other approaches, such as using Gaussian processes (Wilkins 2019), can similarly calculate frequency-resolved lags.

Here, we begin by presenting simulations to demonstrate the ability of frequency-resolved analysis to identify reprocessing from an extended region, before we apply the technique of Zoghbi et al. (2013) to analyze the light curves from the AGN Space Telescope and Optical Reverberation Mapping (STORM) campaign on NGC 5548, and calculate frequency-resolved lags and discuss their implications.

2. Frequency-resolved Lags

In a linearized-reverberation-mapping model, we can relate the reverberating (reprocessed) light curve, $r(t)$, to the driving light curve, $d(t)$, via

$$r(t) = \int_{-\infty}^{\infty} \psi(\tau) d(t - \tau)\, d\tau, \qquad (1)$$

where $\psi(\tau)$ is the impulse response function (often called the transfer function). This response function encompasses the information about the geometry of the reprocessor, and gives the response of the gas to a delta-function flare in the driving continuum. The reprocessed light curve is a blurred and shifted (lagged) version of the driving light curve. Typical time-domain analysis performed in optical AGN studies calculate the cross-correlation function (CCF), which is a convolution of the response function with the auto-correlation function (ACF) of the driving light curve. Thus measuring the centroid of the CCF approximately measures the centroid of the response

function. If the reprocessing region is complex, this will be reflected in the shape of the CCF. For instance, a skewed CCF whose peak and centroid do not align might be indicative of an extended response function. But, most CCF analysis does not attempt to study these effects and stops at measuring the lag via either the peak and/or centroid.

In the Fourier domain (indicated below using upper-case symbols), the Fourier transform of $r(t)$ is simply the Fourier transform of $d(t)$ multiplied by the Fourier transform of the response, i.e.,

$$R(\nu) = \Psi(\nu)D(\nu), \qquad (2)$$

where ν is the Fourier frequency. We can therefore also relate the power spectral density (PSD) of the reprocessed light curve to the driving light curve (see Section 2.4 in Uttley et al. 2014 and also Papadakis et al. 2016), with the reprocessed PSD equal to the PSD of the driving light curve multiplied by the modulus-squared of the Fourier transform of the response function:

$$|R(\nu)|^2 = \Psi^*(\nu)D^*(\nu)\Psi(\nu)D(\nu) = |\Psi(\nu)|^2|D(\nu)|^2, \qquad (3)$$

where * indicates a complex conjugate. To calculate the lag between two light curves in the Fourier domain we use the cross-spectrum, $C(\nu)$, which is the complex conjugate of the Fourier transform of one light curve multiplied by the Fourier transform of the other. Thus,

$$C(\nu) = D^*(\nu)R(\nu) = D^*(\nu)D(\nu)\Psi(\nu) = |(D(\nu)|^2\Psi(\nu). \qquad (4)$$

In other words, the cross-spectrum is equal to the driving signal PSD multiplied by the Fourier transform of the response function. The phase of the cross-spectrum gives the phase lag between the light curves at each Fourier frequency (see Sections 2.1.2 and 2.4.3 of Uttley et al. 2014, for a more detailed description). This phase lag can be converted to a time lag by dividing by $2\pi\nu$. The cross-spectrum can therefore be used to determine lag as a function of frequency, providing more detail about the response function than simply measuring the average value, and allowing an investigation of the timescale (frequency) dependence of the response.

Reverberation acts to smooth light curves and, as described above, therefore changes the light curve's PSD. The normalization of the PSD should drop, and its slope should steepen with the amplitude reduced most at the highest frequencies. Qualitatively this can be easily seen looking at the light curves, with longer-wavelength light curves (which are thought to originate from further out in the accretion disk) being smoother and having lower-variability amplitudes (for one example, see Figure 2 in Fausnaugh et al. 2016). A quantitative analysis by Panagiotou et al. (2020) measures the PSD of these NGC 5548 light curves, finding that they are all well described by a simple power law. The slope of the UV/optical PSDs are the same, but steeper than the X-ray PSD. The normalization of the PSD decreases with wavelength, consistent with thermal reverberation. However, there has not been a study of the cross-spectrum and frequency-resolved lags in AGN light curves across the full UV/optical range. The only previous attempt compared the X-ray and UV light curves in Mrk 335, and only saw tentative lags on long timescales, with all values consistent with zero within 2σ (Griffiths et al. 2021).
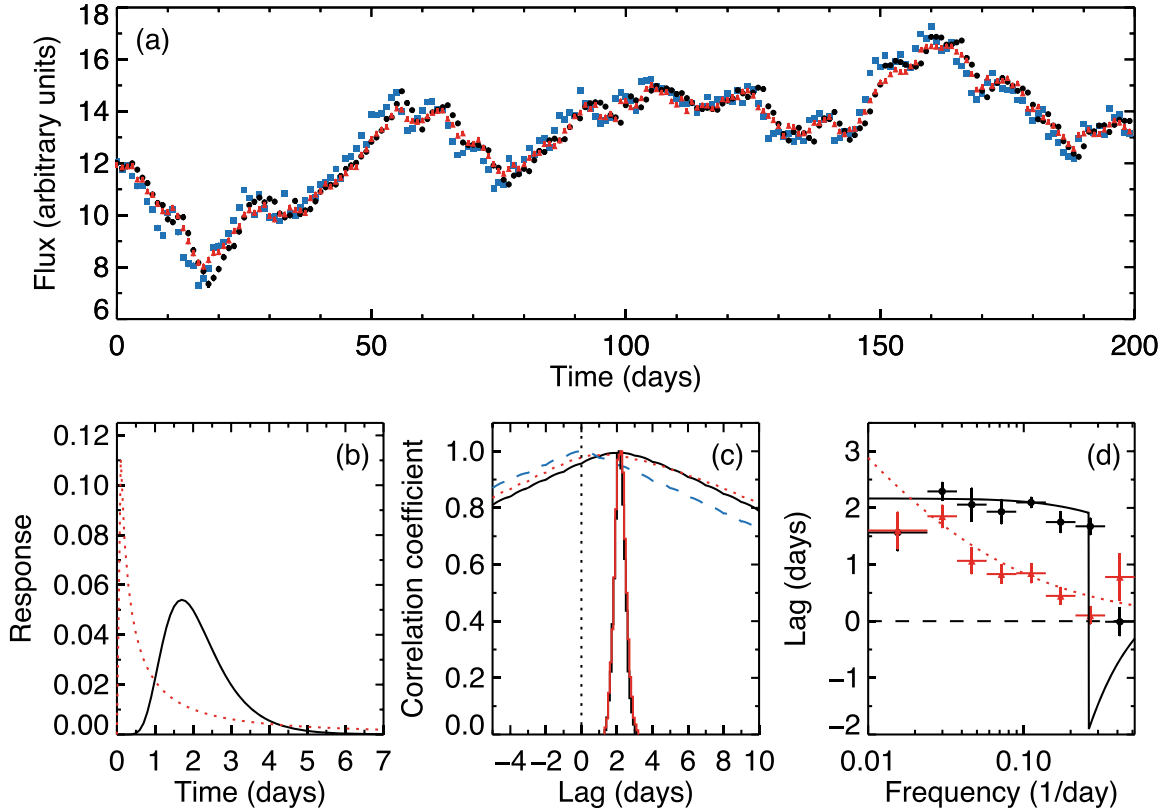
Figure 1. (a) Simulated driving light curve (blue squares), and reprocessed light curves for response 1 (black circles) and response 2 (red triangles). (b) Response function 1 (black solid line) and response function 2 (red dotted line). Response 2 peaks at shorter times, but has a more prominent tail at late times. (c) Auto-correlation function of the driving light curve (blue dashed line) compared to the cross-correlation functions for response 1 (black line) and response 2 (red dotted line). The black and red histograms show the centroid lag distributions from the FR/RSS procedure used to determine the lag uncertainties. The centroid lags for response 1 and 2 are consistent within 1σ. (d) The frequency-resolved lags recovered from the light curves for response 1 (black circles) and response 2 (red triangles), with the model frequency-resolved lags shown as black solid and red dotted lines for response 1 and 2 respectively.

3. Simulations

To demonstrate the potential advantages of undertaking a frequency-resolved lag analysis, we perform some simple simulations. We generate a driving light curve several times longer than needed (to account for red noise leakage; Uttley et al. 2002), using the Timmer & Koenig (1995) algorithm and assuming a PSD with slope −2. We then convolve this with a response function. After this we cut the light curve down to 200 days, and sample with a cadence of once per day, adding random Gaussian noise of 1% (panel (a) of Figure 1 shows the light curves used). We test a couple of different response functions, chosen so that their recovered centroid lag values are approximately the same; however, their different properties mean that one has a more prominent tail at long delays. Both responses are chosen to be log-normal functions:

$$\psi(t) = \frac{1}{S\sqrt{2\pi}\,t} \exp\left[-\frac{(\ln t - M)^2}{2S^2} \right] \tag{5}$$

where t is time in days. For response 1 we choose $M = \ln 2$ and $S = 0.4$. For the second response (response 2) we also use a log-normal distribution, this time with $M = 0$ and $S = 2.0$. The smaller M means that the response peaks at shorter lags, but the larger S broadens the distribution, giving a significant tail to long lags (as can be seen in panel (b) of Figure 1). The peak and centroid of the first response are quite similar with 1.70 and 2.17 days, respectively; the second response, however, has a

peak of 0.02 days and a centroid of 6.85 days. The responses are calculated over the range $t = 0$–1000 days.

With these simulated light curves we test how well distinguished the two response functions are by the standard CCF and frequency-resolved lags approaches. First, we calculate the CCF with respect to the driving light curve and use the standard flux randomization, random subset sampling method (FR/RSS; White & Peterson 1994; Peterson et al. 2004) to determine the peak and centroid lags and their uncertainties. Panel (c) of Figure 1 shows the CCFs and histograms of centroid values from the FR/RSS technique. By eye the CCFs look quite similar, though it can be seen that response 2 peaks at slightly smaller lags and has a higher correlation coefficient at larger lags compared to response 1. Quantitatively, the recovered centroid lags are consistent within their 1σ uncertainties, but the peak lags differ. The recovered peak and centroid lags are $\tau_{\mathrm{peak},1} = 2.20 \pm 0.10$ days, $\tau_{\mathrm{cent},1} = 2.16 \pm 0.29$ days for response 1 and $\tau_{\mathrm{peak},2} = 1.35 \pm 0.10$ days, $\tau_{\mathrm{cent},2} = 2.19 \pm 0.31$ days for response 2. The peak and centroid for response 2 are significantly different from each other, indicating the response function is asymmetric with a tail to long lags. However, in most reverberation analyses the peak and centroid lags are usually simply quoted and rarely are the differences in the peak and centroid lags or the shape of the CCF investigated in more detail.

The differences between the two responses becomes clearer when looking at the frequency-resolved lags. Since the simulations use a strict 1 day cadence, the Fourier analysis

can be performed using standard fast Fourier transform techniques (Uttley et al. 2014) to calculate the cross-spectrum (which is then binned in frequency), and from the phase of the (binned) cross-spectrum calculate the time lag. The resultant lags as a function of Fourier frequency are shown in panel (d) of Figure 1, where we also show the predicted lags based on the model response functions. Here, a clear difference between the frequency-resolved lags is easily seen.

There are notable features in the lag-frequency plot worth describing. Phase wrapping (the sudden flip in sign of the lag) occurs as the phase lags are limited to the range $-\pi$ to π. So, a positive or negative shift of half a wave cannot be distinguished. For a symmetric response function this will occur at a frequency of $\nu = 1/2\tau_0$ (see Section 4.1.1 of Uttley et al. 2014), where τ_0 is the centroid of the response function. For an asymmetric response function τ_0 is skewed to slightly longer than the peak, and, of course, depends on the Fourier transform of the response function. This phase wrapping therefore indicates the approximate lag where the response function peaks. Phase wrapping is clearly seen in response 1, at a frequency of approximately 0.25 day^{-1}, corresponding $\tau_0 = 2$ days. For the second response the phase wrapping occurs at a frequency higher than probed by the cadence of the data sampling.

Aside from the phase wrapping the evolution in lag with frequency is also significantly different for the different models. For response 1 at frequencies lower than where phase wrapping occurs the lag remains approximately constant. This flattening indicates that there is not significant additional response on long timescales. For response 2, however, the short peak and long tail to the response causes the lags to smoothly decrease with increasing frequency; in other words, as you go to greater distances (longer lags), there is additional reprocessing.

At the lowest frequencies once you get beyond the equivalent timescale where the response function has dropped to zero, the lag becomes a constant and corresponds to the centroid lag of the response function. This is clearly seen in panel (d) of Figure 1 for response 1 (black line). However, for response 2 (red dotted line) the 200 day campaign is not long enough to get frequencies low enough to completely reach the flat part of the lag-frequency evolution. The length of the campaign obviously limits the lowest-frequency bins that can be analyzed, and biases can be introduced if there are not enough Fourier frequencies per bin (see e.g., Epitropakis & Papadakis 2016). If there is response from a distant reprocessor, a 200 day campaign is not necessarily long enough to fully recover the lags. This can only be mitigated by having significantly longer monitoring campaigns, which presents a challenge given ground-based visibility of objects.

In summary, while with standard CCF approaches the recovered centroid lags are consistent between the two different responses used in the simulations, when looking at frequency-resolved lags the evolution of the lag with frequency allows one to clearly distinguish between the two responses. This demonstrates the potential in using the technique to better constrain the shape of the response function in UV/optical continuum reverberation and help determine the contribution to the lags from the accretion disk, BLR, or other location.

4. Frequency-resolved Lags in NGC 5548

The simulations demonstrate the potential power in frequency-resolved analysis to distinguish between different shaped response functions. We therefore apply the frequency-resolved lags method

to a real data set. Over the last 5 yr or so a number of intensive, multiband continuum-reverberation-mapping campaigns have taken place using Swift to anchor the analysis, with high cadence (usually multiple times per day) X-ray and UV/optical light curves often spanning several hundred days. AGN STORM (De Rosa et al. 2015) was the first campaign to combine Swift, HST, and ground-based monitoring, and clearly detects wavelength-dependent lags (Edelson et al. 2015; Fausnaugh et al. 2016). With 262 epochs over 169 days in the Swift UVW2 light curve and 171 epochs over 175 days in the HST light curves, it remains one of the best data sets currently available. We therefore use these data to test this analysis method. Since these real data are not evenly sampled in time we cannot use standard fast Fourier transforms, and so instead use a maximum-likelihood method.

The lags are calculated using the maximum-likelihood method outlined in Zoghbi et al. (2013), which was first presented in Miller et al. (2010). The codes used for the analysis can be found at https://zenodo.org/record/5566974 (Zoghbi et al. 2021). A likelihood function for the observed light curves is constructed assuming that the power spectra and the response function are piecewise functions of Fourier frequency, and whose parameters are estimated by numerically maximizing the likelihood. A Gaussian likelihood is used, with a covariance matrix that is built from the Fourier transform of the piecewise functions.

In other words, the autocovariance function is defined as

$$\mathcal{A}(\tau) = \int |D(\nu)|^2 \cos(2\pi\nu\tau)d\nu, \qquad (6)$$

where ν is the Fourier frequency, τ is the time difference between pairs of time points, and $|D(\nu)|^2$ is the power spectral density, modeled here as a piecewise sum over n_ν frequency bins: $|D(\nu)|^2 = \sum_i D_i$. Likewise, the cross-covariance is defined as

$$\mathcal{X}(\tau) = \int |D(\nu)|^2 |\Psi(\nu)| \cos(2\pi\nu\tau - \phi(\nu))d\nu, \qquad (7)$$

where $|\Psi(\nu)|$ and $\phi(\nu)$ are the amplitude and phase of the response function. These two are also modeled as piecewise functions of frequency ($\Psi(\nu) = \sum_i \Psi_i$ and $\phi(\nu) = \sum_i \phi_i$). Assuming the observed light curves are generated by such general Gaussian processes, a likelihood function can be constructed and maximized to obtain best values Ψ_i and ϕ_i as a function of frequency. The time lag is then obtained by dividing the phase ϕ_i by $2\pi\nu_i$. All the lags are measured relative to the light curve at 1158 Å, which we take as the reference.

We report lags in six logarithmically spaced frequency bins between 0.012 and 1.11 days^{-1} in Table 1. Two additional frequency bins at the start and end of this range (0.0017–0.012 and 1.11–2700 days^{-1}) were included in the maximum-likelihood calculations, but were not in the modeling because not all light curves have information in this range, and also because the first and last bin tend to show biases in their measured lags (Zoghbi et al. 2021; see also Epitropakis & Papadakis 2016).

The uncertainties in the reported parameters are obtained by running Monte Carlo Markov Chains starting with a random set of parameters around the best fit obtained by numerical maximization of the likelihood function. We use the 1σ uncertainty of the lag as the standard deviation of the parameter chain. To allow subsequent modeling, we ensured that all measured lags have Gaussian probability densities. If any of the

Table 1
Lags (Measured with Respect to 1158 Å) in Each Fourier Frequency Bin

Filter/ Wavelength	Lag (days) in Fourier Frequency Range				
	0.012–0.025 day^{-1}	0.025–0.054 day^{-1}	0.054–0.115 day^{-1}	0.115–0.244 day^{-1}	0.244–1.111 day^{-1}
1367 Å	0.59 ± 0.19	0.20 ± 0.10	0.06 ± 0.10	0.10 ± 0.08	−0.03 ± 0.38
1479 Å	0.17 ± 0.27	0.42 ± 0.11	0.17 ± 0.11	0.09 ± 0.09	0.04 ± 0.14
1746 Å	0.77 ± 0.28	0.37 ± 0.14	0.33 ± 0.13	0.25 ± 0.12	0.00 ± 0.31
UVW2	1.41 ± 0.39	0.10 ± 0.20	−0.12 ± 0.18	−0.23 ± 0.14	0.01 ± 0.48
UVM2	1.47 ± 0.59	−0.02 ± 0.22	−0.03 ± 0.23	−0.18 ± 0.16	0.05 ± 0.46
UVW1	1.41 ± 0.58	0.32 ± 0.26	0.11 ± 0.22	−0.13 ± 0.20	0.05 ± 0.45
U$_{Swift}$	2.54 ± 0.68	1.08 ± 0.31	0.36 ± 0.39	0.24 ± 0.44	−0.05 ± 0.50
u	4.12 ± 0.51	1.66 ± 0.38	1.65 ± 0.49	0.57 ± 0.46	0.12 ± 0.25
B	3.85 ± 0.46	1.79 ± 0.21	1.09 ± 0.42	0.19 ± 0.16	
B$_{Swift}$	1.94 ± 0.88	0.89 ± 0.44	0.72 ± 1.69	0.10 ± 0.66	
g	3.48 ± 0.57	1.73 ± 0.34	1.61 ± 0.31	0.70 ± 0.37	0.02 ± 0.27
V	3.74 ± 0.39	2.72 ± 0.20	0.95 ± 0.22	0.65 ± 0.28	0.04 ± 0.47
r	5.74 ± 0.51	2.75 ± 0.38	2.62 ± 1.29	−0.18 ± 1.45	0.22 ± 0.41
R	4.25 ± 0.71	3.13 ± 0.37	1.55 ± 1.01	1.33 ± 0.84	−0.02 ± 0.48
i	5.68 ± 0.55	3.64 ± 0.33	2.55 ± 0.40	0.36 ± 0.39	
I	5.57 ± 0.67	3.15 ± 0.41	0.42 ± 3.07	0.31 ± 0.26	
z	4.96 ± 0.65	4.07 ± 0.35	3.09 ± 0.50	−0.10 ± 1.49	−0.03 ± 0.17

frequency bins is not consistent with a Gaussian distribution, we group it with a neighboring frequency to increase the signal and ensure normal uncertainties. As a consequence, we combined the two highest frequency bins in all light curves (resulting in five frequency bins), and for the lower cadence and/or signal light curves we end up with four frequency bins. The recovered frequency-resolved lags are shown in Figure 2.

For comparison we also show the cross-correlation lags of Fausnaugh et al. (2016) in Figure 2 as a blue horizontal region. The cross-correlation lags are most consistent with frequency-resolved lags at frequencies lower than 0.1 day^{-1}, with the frequency bin from 0.025–0.054 day^{-1} often consistent with the cross-correlation lag. In Figure 3 we show the lags as a function of wavelength in each of the frequency bins. For wave bands with five frequency bins we average the lags between the two highest frequency bins in order to be able to compare with the wave bands with four frequency bins. The lags in the three lowest frequency bins all approximately follow $\tau \propto \lambda^{4/3}$ (dashed line), while the highest frequency bin is dramatically different, with most lags consistent with zero.

The best-fitting $\tau = \tau_0(\lambda/\lambda_0 - 1)^{4/3}$ from the cross-correlation analysis of Fausnaugh et al. (2016) gives $\tau_0 = 0.42 \pm 0.02$ days for their reference band of 1367 Å. Adjusting to 1158 Å by subtracting the 1158 to 1367 Å lag of 0.17 ± 0.16 days gives $\tau_0 = 0.25 \pm 0.16$ days, consistent with the $\tau_0 = 0.32 \pm 0.01$ days we get from fitting the lags in the 0.025–0.054 day^{-1} frequency bin. This further demonstrates how the CCF analysis is picking up the lags at low frequencies, corresponding approximately to size scales greater than 18 light days. Given the red-noise nature of AGN variability, the light curves have the most power at the lowest frequencies. Filtering out the long timescale variability in the light curves (i.e., by subtracting a running box-car average) leads to shorter CCF lags.

5. Modeling

A clear trend is seen when looking at the lag-frequency plots for each wave band shown in Figure 2. Generally, they show a smoothly decreasing lag with increasing frequency. The lag is consistently longest at the lowest frequency examined, and

close to zero above a frequency of 0.1 day^{-1} in most bands. Comparing to the simulations presented in Section 3, we can immediately rule out a close-to-symmetric response function. Such a response function should show an almost flat lag-frequency spectrum up to the frequency where phase wrapping occurs. For instance, a centroid lag of 2 days (e.g., similar to the g-band lag) should show phase wrapping at 0.25 day^{-1}. Since 0.25 day^{-1} is a higher frequency than probed here we would expect a lag of 2 days in all frequency bins. The simulations also show that a smooth decrease in lag with increasing frequency results from an asymmetric response function with a long tail at larger lags. This immediately suggests that the continuum lags we observe come from an extended reprocessing region. The fact that the lags are very short at frequencies greater than 0.1 day^{-1} suggests reprocessing on size scales less than 10 light days happens very rapidly.

We note that accretion-disk response functions are typically not symmetric and do show a tail at long lags. For instance, see examples in Cackett et al. (2007), Starkey et al. (2016) and Kammoun et al. (2019). Thus, to test whether accretion-disk reprocessing models show the observed lag-frequency spectrum, we use the models described in detail in Cackett et al. (2007). Those models are characterized by the disk temperature in the bright and faint state, T_B and T_F. We set T_B and T_F so that the centroid lag approximately follows the observed CCF lag centroids from Fausnaugh et al. (2016), getting a good match with $T_B = 20,900$ K and $T_F = 16,800$ K. We assume an inclination of 45°.

To calculate the model lags we note that we have calculated the observed lags with respect to the 1158 Å band. In the disk-reprocessing model this is not the driving light curve itself, and is itself a reprocessed light curve. To correct for this we note that if the light curve in band A is $a(t)$ and has response function $\psi_A(t)$, and the light curve in band B is $b(t)$ and has response $\psi_B(t)$, then the Fourier transform of each light curve is

$$A(\nu) = \Psi_A(\nu)D(\nu), \qquad (8)$$
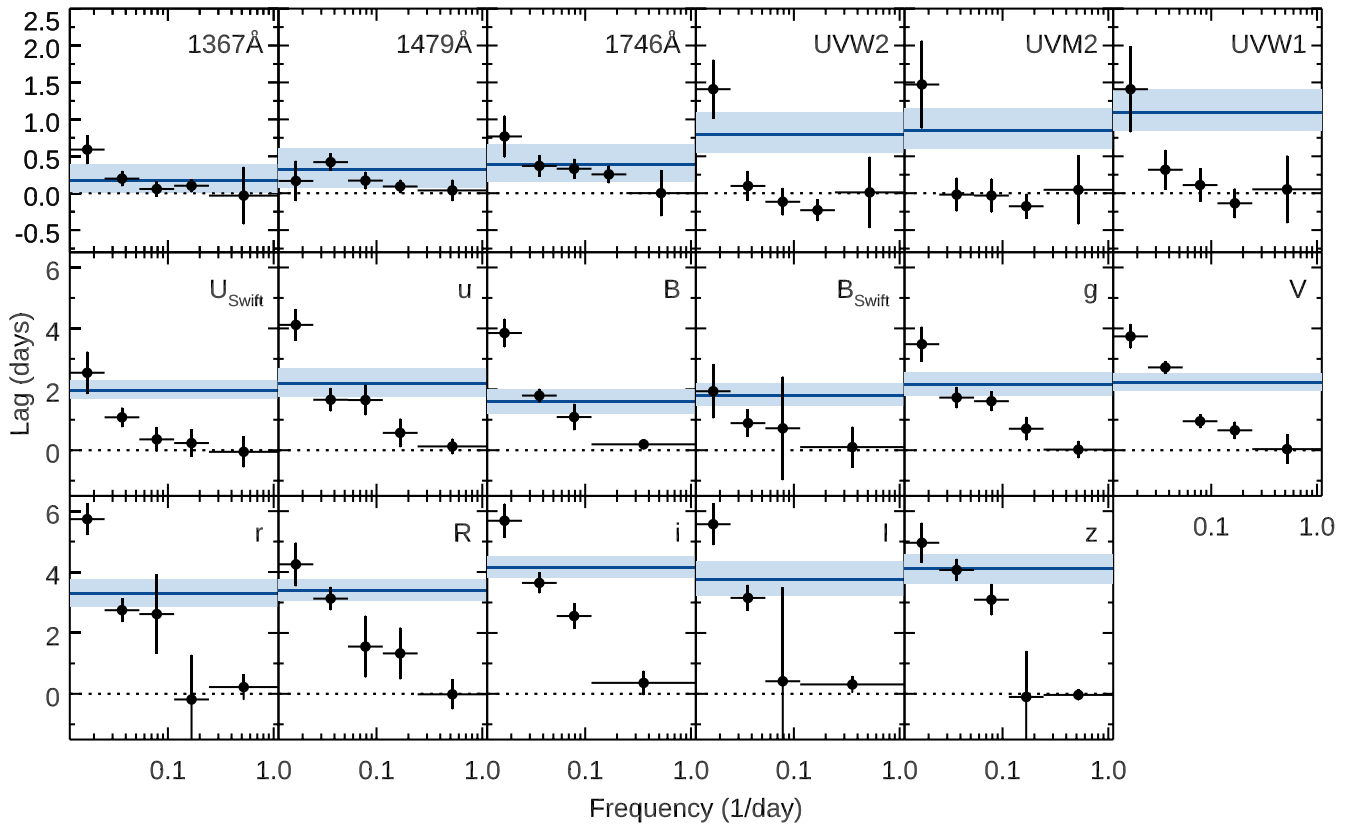
$$B(\nu) = \Psi_B(\nu)D(\nu). \qquad (9)$$

Figure 2. Frequency-resolved lags in NGC 5548. Blue lines are the cross-correlation function lags from Fausnaugh et al. (2016) for comparison. Note that the *y*-axis range in the top panel is narrower than for the other two panels to better visualize the shorter lags in those bands.

The cross-spectrum of light curve B with respect to A will then be

$$C(\nu) = A^*(\nu)B(\nu) = |D(\nu)|^2\Psi_A^*(\nu)\Psi_B(\nu). \qquad (10)$$

Since $|D(\nu)|^2$ is a set of real constants (and therefore does not affect the complex phase), the lags can be calculated from the phase of $\Psi_A^*(\nu)\Psi_B(\nu)$ alone. Here, we adopt the 1158 Å response as $\psi_A(t)$, and the response for each band in turn is $\psi_B(t)$. Since the lag between the driving light curve and 1158 Å is very short in these models (approximately 0.2 days), this is only a minor adjustment to the lags. It would be significantly more important if an optical band was used as the reference.

The resulting lags from this disk model are shown as blue lines in Figure 4. While the model does well at matching the lag-frequency spectrum in the *I* and *z* bands, at shorter wavelengths the disk model consistently underpredicts the lag at the lowest frequency and at the shorter wavelengths overpredicts the lag at the highest frequencies. In all but the longest wavelength bands, the disk model starts to flatten and only decreases slightly in lag below frequencies of around 0.1 day^{-1}.

In order to get longer lags at low frequencies and shorter lags at higher frequencies, we test a model for two reprocessors with reverberation from a significantly smaller disk and reverberation from a more distant, BLR-like reprocessor. The best-fitting lag–wavelength relation from Fausnaugh et al. (2016) had a normalization about a factor of 3 larger than expected for reasonable assumptions about mass and accretion rate. Our first disk model test shows that the lag-frequency spectrum is not consistent with this. For this second model we therefore adopt a

disk model that is the size expected for a standard disk. Fausnaugh et al. (2016) state such a disk model would have a normalization leading to a lag at the reference band of 0.14 days. We adjust T_B and T_F to match this, finding $T_B = 6400$ and $T_F = 5000$ K. Again, we assume an inclination of 45°. In addition to this disk model, we add a log-normal shaped response (Equation (5)) to account for continuum emission from the BLR. During the AGN STORM campaign, the Hβ lag was measured to be approximately 6 days when using the 1158 Å reference band (Pei et al. 2017). Since the median of a log-normal distribution is e^M, we therefore assume $M = \ln 6$, and to give reprocessing over an extended region assume $S = 1.0$. We note these are just adopted for demonstrative purposes. These two responses are then added together as follows:

$$\psi_2(t) = (1 - f)\psi_{\text{disk}}(t) + f\psi_{\text{BLR}}(t), \qquad (11)$$

where $\psi_{\text{disk}}(t)$ is the disk response and $\psi_{\text{BLR}}(t)$ the BLR response, which are both normalized to have a total area of 1. f is then the response fraction from the BLR and $\psi_2(t)$ is the combined response. We show three examples of this combined response function in Figure 5. When $f = 0.1$ (blue solid line) the response is dominated by the accretion-disk component, and has a prominent peak very close to zero. At the other extreme, when $f = 0.9$ (green dotted line) the response is dominated by the extended component and has a more prominent tail at long lags.

We apply this model to the data, fitting for f in each band. The resulting lag-frequency models are shown as green lines in Figure 4. This second model does a much better job of matching the observed frequency-resolved lags. Quantitatively,
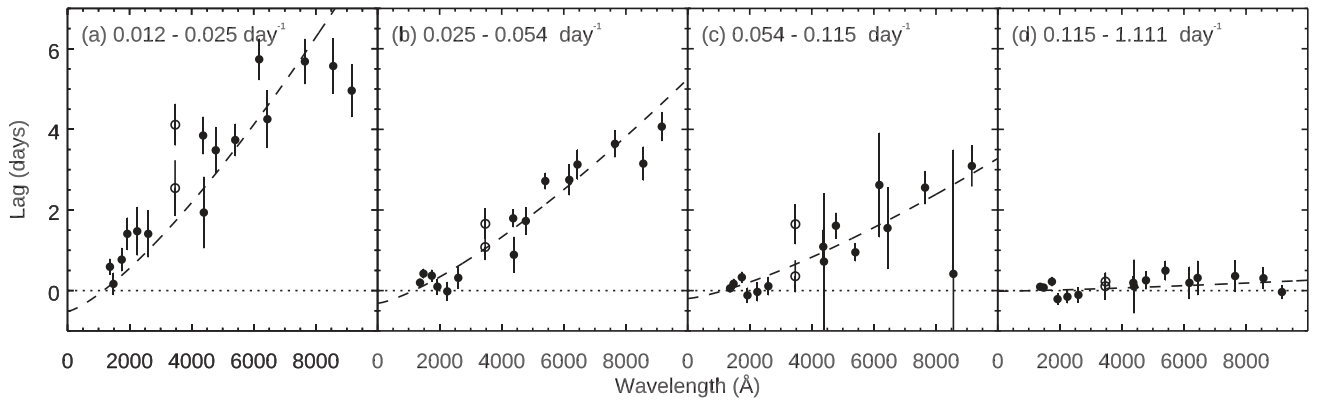
Figure 3. Lags as a function of wavelength for each of the frequency bins. The dashed line is the best-fitting $\tau \propto \lambda^{4/3}$ relation. The Swift U and ground-based u bands are shown as open circles.

the χ^2 for the large-disk model = 236.4 for 81 data points, giving $\chi^2_\nu = 2.92$. For the small-disk-plus-BLR model we get $\chi^2 = 70.95$ for 81 data points and 17 free parameters (f in each band), giving $\chi^2_\nu = 1.11$, i.e., a significantly improved fit compared to the large-disk model. The second model does better since the short disk lags give short lags above a frequency of 0.1 day^{-1}, while the BLR response allows the lags to match at the lowest frequencies and gives a smooth decrease in lag with increasing frequency. A breakdown of the model components is shown for the V band in Figure 6. The disk model (red line) provides a short ($\sim$0.25 day) lag at all observed frequencies, while the BLR model (blue line) shows decreasing lag with increasing frequency up to where phase wrapping takes place. The combination of the two with $f = 0.63$ matches the observed lags.

Since the disk lags are short ($<$1.5 day in all bands), the BLR fraction has to increase significantly with increasing wavelength, which we show in Figure 7. This shows a generally increasing fraction with increasing wavelength. The u band is offset from the general trend, as would be expected if this distant reprocessor component is associated with diffuse continuum from the BLR (Korista & Goad 2001), though we note the Swift U does not show a significant offset. The u/U bands are shown as open circles in Figures 3 and 7.

For this two-component model we also tried allowing M and S to be free parameters in the fit, but common for all wave bands. We explore the range for e^M from 0.5–11.5 days. and S from 0.1 to 5.5. We find a best fit with parameters (and 1σ uncertainties): $e^M = 7.3^{+0.7}_{-0.9}$ days and $S = 1.1 \pm 0.2$; this gives $\chi^2 = 68.5$ for 62 degrees of freedom. There are strong constraints on the lower limit for e^M, with a 3σ lower confidence limit of 5.0 days. The upper range is not as tightly constrained, with the best fit at $e^M = 11.5$ days at approximately the 2σ confidence level. Again, this indicates that this more distant component is consistent with the size scale of the BLR in NGC 5548.

We also tested fitting the lag-frequency spectra with a single log-normal component, allowing M and S as free parameters in each wave band. While this does give good fits at each wave band, we find that in the UV bands M tends to be as small as allowed and S wants to be large (in the range 4–5 for the shortest wavelength bands). This gives a response function that peaks as close to zero as possible yet has a long tail. This is needed to have the lags close to zero at most frequencies but still give a nonzero lag at the lowest frequencies. This is

equivalent to our two-component model described above with a small disk and larger reprocessor.

To summarize, the modeling we have performed demonstrates that the frequency-resolved continuum lags in NGC 5548 are more consistent with a standard-sized accretion disk with an additional contribution from BLR continuum emission than emission from a larger than expected accretion disk alone.

6. Discussion

We have introduced the use of frequency-resolved lag analysis in UV/optical continuum-reverberation mapping. Since this method determines the lag at each Fourier frequency, it separates out the timescale (and hence the corresponding size scale) at which the light curves are lagging each other. While this information is also encoded in the auto- and cross-correlation functions, it is not easy to assess with standard cross-correlation analysis.

To demonstrate the usefulness of frequency-resolved lag analysis, we simulated light curves assuming two different reprocessing geometries: one that has a response function that is close to symmetrical, and another that has a peak at shorter lags but a significant tail to longer lags. Standard cross-correlation lag analysis recovers centroid lags for the two different light curves that are consistent with each other. However, when using frequency-resolved lag analysis they can be easily distinguished.

We apply frequency-resolved lag analysis to the UV/optical continuum light curves of NGC 5548 from the AGN STORM campaign (Edelson et al. 2015; Fausnaugh et al. 2016). This campaign spanned approximately 170 days with daily monitoring by the HST and an average of 1.5 observations per day with Swift, as well as ground-based monitoring, leading to light curves spanning from 1158 Å to z. Since the light curves are not strictly evenly sampled and do contain occasional gaps we applied the maximum-likelihood method of Zoghbi et al. (2013), resulting in lag measurements in four or five frequency bins (depending on the quality of the light curve and sampling) over the frequency range 0.012 to 1.11 day^{-1}. The lags are consistently longest in the lowest-frequency bin and through the optical they decrease with increasing frequency. In the UV many of the lags outside the lowest frequency are consistent with zero. This is similar to the findings of Griffiths et al. (2021), who look at the lag between X-rays and UVW2 in Mrk 335, only finding tentative evidence for nonzero lags in the lowest-frequency bins.
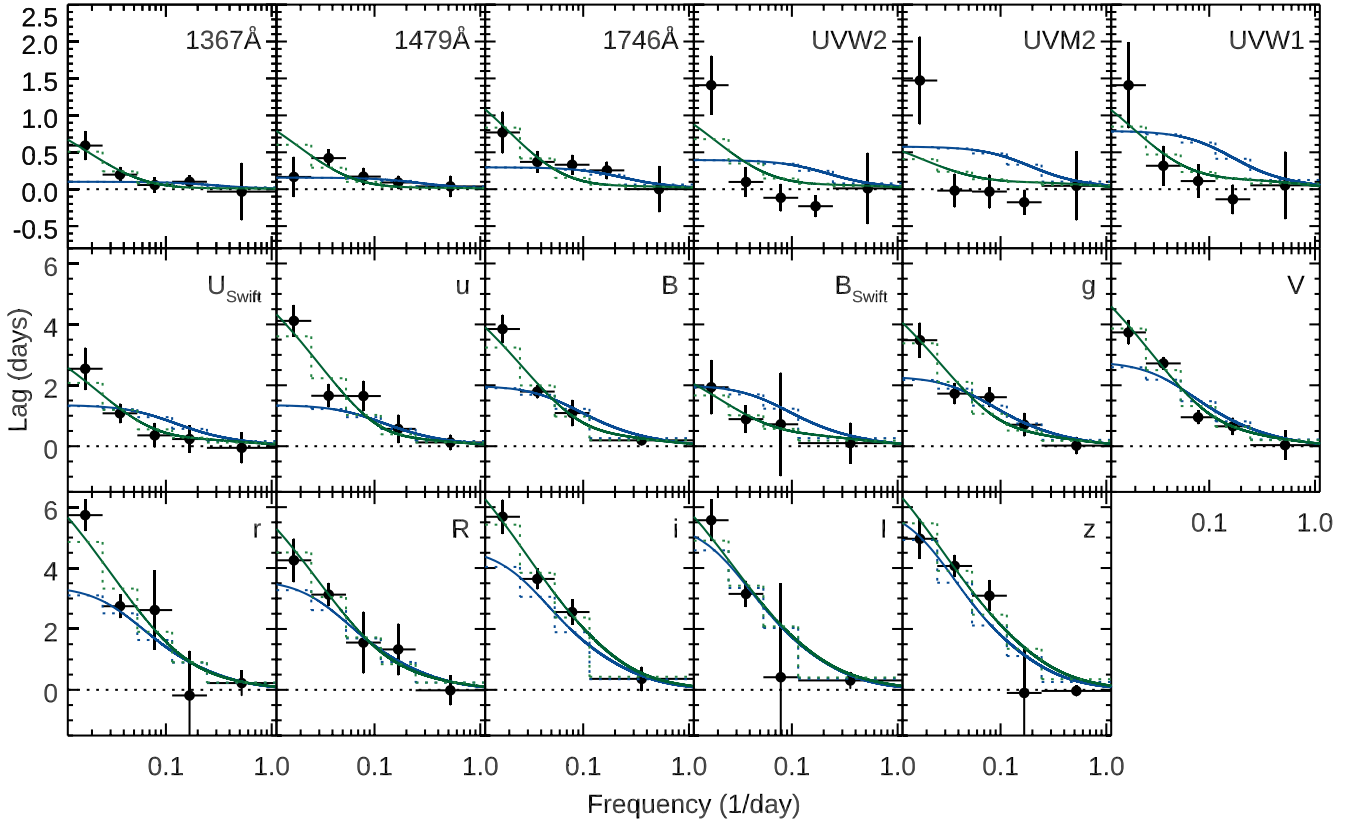
Figure 4. Frequency-resolved lags in NGC 5548 with different reprocessing models. The blue lines show reprocessing from an accretion disk only while the green lines show a model consisting of a combination of significantly smaller disk and an extended reprocessor (dotted lines show the model averaged over each frequency bin); see the text for more details.
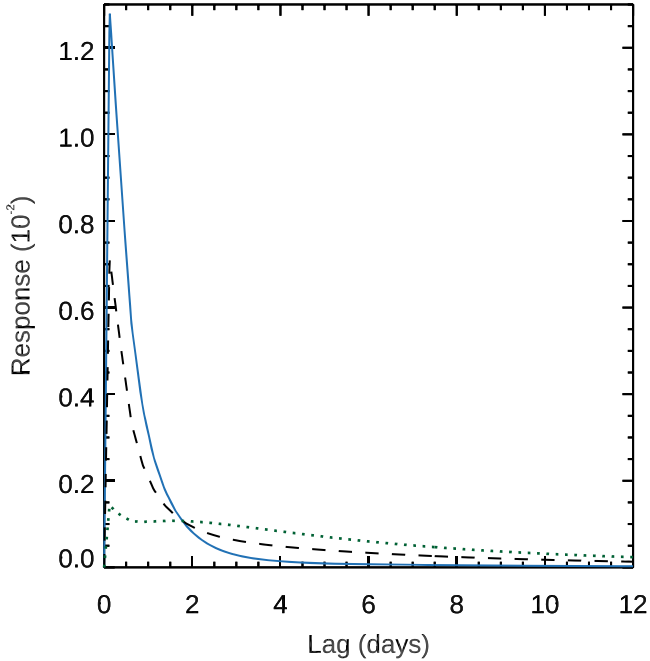

Figure 5. Examples of the combined disk and BLR model response function for different values of the BLR fraction, f. We use the disk response for the V band combined with the BLR model with $f = 0.1$ (blue solid line), $f = 0.5$ (black dashed line), and $f = 0.9$ (green dotted line).
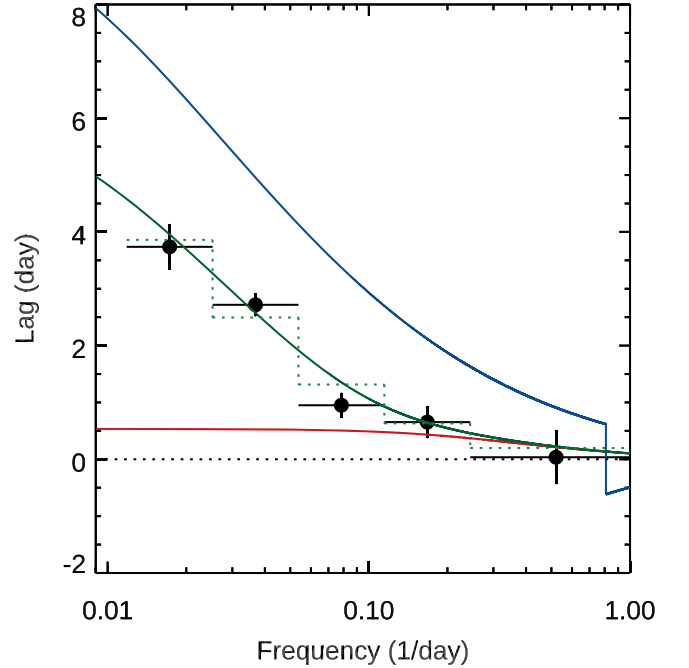

Figure 6. Frequency-resolved lags in the V band. The green line shows the best-fitting disk+BLR model. The individual components are shown as red (disk) and blue (BLR) lines.

The lags do not evolve dramatically over the frequency range 0.1–1 day^{-1}, which indicates that reprocessing on size scales less than 10 light days is occurring promptly. The lag

signal detected by CCF analysis corresponds to continuum emission happening on size scales greater than this. The lags in the frequency bin centered at 0.037 day^{-1} approximately
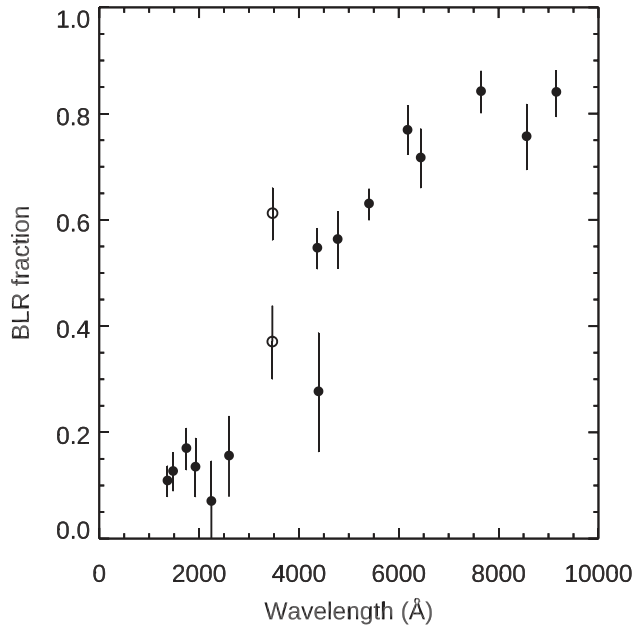
Figure 7. Fraction of the response from the BLR component from the two-component (small disk plus BLR) model. The Swift *U* and ground-based *u* bands are shown as open circles.

follow a $\tau \propto \lambda^{4/3}$ relation that is consistent with the lags recovered from cross-correlation analysis by Fausnaugh et al. (2016). Thus, the cross-correlation analysis is picking out the lags occurring on timescales of around 27 days.

We use the accretion-disk response functions of Cackett et al. (2007) calculated to match the observed $\tau \propto \lambda^{4/3}$ relation from cross-correlation analysis. The frequency-resolved lags from these response functions do not match the observed frequency-resolved lags well through the UV and much of the optical, though do better at the longest wavelengths. They show lags that flatten off at the lowest frequencies and generally fail to match the lag in the lowest-frequency bin. This suggests the need for an additional reprocessor at large distances.

As an alternative model, we try a significantly smaller disk (consistent with expections for a Shakura & Sunyaev 1973 disk), with an additional extended reprocessor that is a log-normal distribution with median lag of 6 days, consistent with the location of the BLR during the AGN STORM campaign (Pei et al. 2017). We add the responses together but allow the fraction from the BLR component to vary from band to band. We find that this provides a significantly better fit to the frequency-resolved lags, showing smoothly decreasing lag with increasing frequency. This model shows an increasing BLR fraction with increasing wavelength.

We note that it is possible to use a single log-normal response function to fit the lags. However, in order to fit the UV lag-frequency data there needs to be a prompt response plus an extended tail to give zero lag at most frequencies yet still have a nonzero lag in the lowest frequency bin. This is not needed for the longer-wavelength bands, leading to vastly different values for M in each band, which is not easy to interpret physically. On the other hand, the advantage of the two-component disk-plus-BLR model is that the BLR has the same response in each wave band, and the difference in lag measured depends on what fraction of the variable flux comes from the disk (prompt response) versus the BLR (extended response).

We stress that we have not explored the range of possible models in full here, and have only tested a couple of very specific scenarios. The fits presented here do not allow the disk properties to be free parameters; moreover, we use a very simple analytic prescription to represent the BLR. Future work should involve a fuller exploration of disk models; for instance, higher-inclination disks have response functions that peak more prominently at short lags (e.g., see Figure 3 from Starkey et al. 2016). Moreover, an investigation of more complex disk models that fully take into account general relativity (Kammoun et al. 2019, 2021a, 2021b), and including more realistic, physically motivated BLR models (e.g., Korista & Goad 2019; Netzer 2020) is warranted. In particular, the lag models will also need to be consistent with the spectral energy distribution of the variable component as determined by recent flux-flux analysis (e.g., McHardy et al. 2018; Cackett et al. 2020; Hernández Santisteban et al. 2020). Regardless of the details of the model, the frequency-resolved lags demonstrate that the reprocessing gas has both a prompt response (indicated by the short lags at frequencies above 0.1 day^{-1} and the lack of phase wrapping) and an extended response over a range of radii that gives a smooth increase in lag to lower frequencies.

The obvious physical origin for this extended response is diffuse continuum emission from the BLR (Korista & Goad 2001, 2019; Lawther et al. 2018; Netzer 2020, 2022). As described in the Section 1, the excess lag in the u/U bands and around the Balmer jump (e.g., Edelson et al. 2019; Cackett et al. 2018) strongly suggests that there is a significant contribution to the variable flux from the BLR diffuse continuum (however, this excess does not show up as strongly in the frequency-resolved lags here). Moreover, analysis that removes long-term variability via subtracting off a moving box-car average effectively acts as a high-pass filter and finds shorter lags (McHardy et al. 2014, 2018), consistent with what we find here. Further evidence for significant BLR diffuse continuum comes from the decoupling of the continuum variations in NGC 5548 during the AGN STORM campaign (Goad et al. 2019). The emission lines in NGC 5548 were observed to show anomalous behavior where they decorrelated from the UV continuum variations for a period of 60–70 days (Goad et al. 2016). Flux-flux analysis of the continuum variations shows that they also show anomalous behavior during the same period, indicating a significant contribution to the continuum from the BLR (Goad et al. 2019). Netzer (2022) demonstrates that both the shape and normalization of the observed lag spectra are consistent with BLR diffuse emission with a covering factor of ∼0.2.

Our model of a small disk and extended reprocessor whose flux fraction increases with wavelength supports the findings of Chelouche et al. (2019). They perform a bivariate lag analysis of Mrk 279, assuming the variations consist of a driving component that varies nearly simultaneous across the optical and an additional delayed component that partly contributes to each band. They find the fraction of variable flux from the delayed component increases with wavelength and dominates at the longest wavelengths. This is consistent with our best-fitting model where the response from the disk gives a very short lag at all wavelengths, and the fraction of response from the extended region increases with wavelength and is dominant in the reddest bands.

The observed frequency-resolved lag behavior can also be used to test alternative models for the continuum lags. In the

corona-heated accretion-disk reprocessing model (Sun et al. 2020a, 2020b) magnetic fluctuations in the corona affect the disk MHD turbulence and heating rate, driving temperature changes in the disk. The frequency-resolved lags from that model (see Figure 7 of Sun et al. 2020a) decrease smoothly with increasing frequency, broadly consistent with what we observe in NGC 5548.

Over the last five years or so a sample of seven additional AGN have published intensive accretion-disk reverberation-mapping campaigns involving Swift with similarly high-cadence data sets to NGC 5548 (Edelson et al. 2017, 2019; Cackett et al. 2018, 2020; McHardy et al. 2018; Hernández Santisteban et al. 2020; Vincentelli et al. 2021; Kara et al. 2021), and several more are either ongoing or recently completed. In future work we will explore frequency-resolved lags in those data sets to see whether they also display the same-shaped lag-frequency spectra. Furthermore, this technique can be applied to emission-line reverberation too as an additional way to assess the shape of the response function. In this case, it will be important to take account of the response function between the driving light curve and the 5100 Å emission to properly assess the frequency-resolved lags. For continuum-reverberation mapping the fact that the frequency-resolved lags are longest at the lowest frequencies drives the need to want as long a monitoring campaign as possible. Only when the frequency-resolved lags flatten off have we probed the longest timescales where the gas is no longer responding. Similarly, for emission-line reverberation longer campaigns will better capture the full response of the BLR. This should be considered when planning future campaigns.

ORCID iDs

Edward M. Cackett ⓘ https://orcid.org/0000-0002-8294-9281
Abderahmen Zoghbi ⓘ https://orcid.org/0000-0002-0572-9613

References

Cackett, E. M., Bentz, M. C., & Kara, E. 2021, iScience, 24, 102557
Cackett, E. M., Chiang, C.-Y., McHardy, I., et al. 2018, ApJ, 857, 53
Cackett, E. M., Gelbord, J., Li, Y.-R., et al. 2020, ApJ, 896, 1
Cackett, E. M., Horne, K., & Winkler, H. 2007, MNRAS, 380, 669
Chelouche, D., Pozo Nuñez, F., & Kaspi, S. 2019, NatAs, 3, 251
Collier, S. J., Horne, K., Kaspi, S., et al. 1998, ApJ, 500, 162
De Rosa, G., Peterson, B. M., Ely, J., et al. 2015, ApJ, 806, 128
Edelson, R., Gelbord, J., Cackett, E., et al. 2017, ApJ, 840, 41
Edelson, R., Gelbord, J., Cackett, E., et al. 2019, ApJ, 870, 123
Edelson, R., Gelbord, J. M., Horne, K., et al. 2015, ApJ, 806, 129
Epitropakis, A., & Papadakis, I. E. 2016, A&A, 591, A113
Fausnaugh, M. M., Denney, K. D., Barth, A. J., et al. 2016, ApJ, 821, 56
Goad, M. R., Knigge, C., Korista, K. T., et al. 2019, MNRAS, 486, 5362
Goad, M. R., Korista, K. T., De Rosa, G., et al. 2016, ApJ, 824, 11
Griffiths, R.-R., Jiang, J., Buisson, D. J. K., et al. 2021, ApJ, 914, 144
Haardt, F., & Maraschi, L. 1991, ApJL, 380, L51
Hernández Santisteban, J. V., Edelson, R., Horne, K., et al. 2020, MNRAS, 498, 5399
Horne, K. 1994, in ASP Conf. Ser. 69, Reverberation Mapping of the Broad-Line Region in Active Galactic Nuclei, ed. P. M. Gondhalekar, K. Horne, & B. M. Peterson (San Francisco, CA: ASP), 23
Kammoun, E. S., Dovčiak, M., Papadakis, I. E., Caballero-García, M. D., & Karas, V. 2021a, ApJ, 907, 20
Kammoun, E. S., Papadakis, I. E., & Dovčiak, M. 2019, ApJL, 879, L24
Kammoun, E. S., Papadakis, I. E., & Dovčiak, M. 2021b, MNRAS, 503, 4163
Kara, E., Fabian, A. C., Cackett, E. M., et al. 2013, MNRAS, 434, 1129
Kara, E., Mehdipour, M., Kriss, G. A., et al. 2021, ApJ, 922, 151
Korista, K. T., & Goad, M. R. 2001, ApJ, 553, 695
Korista, K. T., & Goad, M. R. 2019, MNRAS, 489, 5284
Lawther, D., Goad, M. R., Korista, K. T., Ulrich, O., & Vestergaard, M. 2018, MNRAS, 481, 533
McHardy, I. M., Cameron, D. T., Dwelly, T., et al. 2014, MNRAS, 444, 1469
McHardy, I. M., Connolly, S. D., Horne, K., et al. 2018, MNRAS, 480, 2881
Miller, L., Turner, T. J., Reeves, J. N., et al. 2010, MNRAS, 403, 196
Netzer, H. 2020, MNRAS, 494, 1611
Netzer, H. 2022, MNRAS, 509, 2637
Pahari, M., McHardy, I. M., Vincentelli, F., et al. 2020, MNRAS, 494, 4057
Panagiotou, C., Papadakis, I. E., Kammoun, E. S., & Dovčiak, M. 2020, MNRAS, 499, 1998
Papadakis, I., Pecháček, T., Dovčiak, M., et al. 2016, A&A, 588, A13
Pei, L., Fausnaugh, M. M., Barth, A. J., et al. 2017, ApJ, 837, 131
Peterson, B. M., Ferrarese, L., Gilbert, K. M., et al. 2004, ApJ, 613, 682
Sergeev, S. G., Doroshenko, V. T., Golubinskiy, Y. V., Merkulova, N. I., & Sergeeva, E. A. 2005, ApJ, 622, 129
Shakura, N. I., & Sunyaev, R. A. 1973, A&A, 24, 337
Starkey, D. A., Horne, K., & Villforth, C. 2016, MNRAS, 456, 1960
Stripe, G. M., Winge, C., Altieri, B., et al. 1994, ApJ, 425, 609
Sun, M., Xue, Y., Brandt, W. N., et al. 2020a, ApJ, 891, 178
Sun, M., Xue, Y., Guo, H., et al. 2020b, ApJ, 902, 7
Timmer, J., & Koenig, M. 1995, A&A, 300, 707
Uttley, P., Cackett, E. M., Fabian, A. C., Kara, E., & Wilkins, D. R. 2014, A&ARv, 22, 72
Uttley, P., McHardy, I. M., & Papadakis, I. E. 2002, MNRAS, 332, 231
Vincentelli, F. M., McHardy, I., Cackett, E. M., et al. 2021, MNRAS, 504, 4337
Wanders, I., Peterson, B. M., Alloin, D., et al. 1997, ApJS, 113, 69
White, R. J., & Peterson, B. M. 1994, PASP, 106, 879
Wilkins, D. R. 2019, MNRAS, 489, 1957
Zoghbi, A., Miller, J. M., & Cackett, E. 2021, ApJ, 912, 42
Zoghbi, A., Reynolds, C., & Cackett, E. M. 2013, ApJ, 777, 24