

X-Ray Reverberation in Active Galactic Nuclei

Eleanor Legg

Mansfield College

University of Oxford

DPhil

2015

Abstract

Narrow Line Type-1 Seyfert active galaxies can exhibit a high degree of variability in the X-ray regime. This thesis examines that variability in the context of reverberation models, in which a flare in activity has an extended, energy dependent, response. A novel method is developed for estimating the response function in different energy bands. This method is then applied to three AGN: Ark 564, 1H 0707–495, and NGC 4051. The striking evidence for reverberation revealed in Ark 564 leads to a more thorough examination of that object, combining spectral and temporal approaches to develop a plausible physical model for its behaviour. The preferred model is one in which the reverberation is due to scattering from hot Comptonizing material approximately 1500 light-seconds from the central source. This conclusion is reinforced by a simulation of the angular dependence of reflection by Comptonizing gas.

Contents

1	Introduction	7
1.1	Active Galactic Nuclei	7
1.2	The Components of the Central Engine of Seyfert Galaxies	10
1.2.1	The Black Hole	10
1.2.2	The Accretion Disk	11
1.2.3	The ‘Corona’	12
1.2.4	Absorbers	13
1.2.5	Wind	13
1.3	Probing the Central Regions	14
1.3.1	Timing Analysis	16
1.3.2	Energy Spectral Analysis	18
1.4	Power Spectra	19
1.5	Lag Spectra	21
1.5.1	Definition	21
1.5.2	History	22
1.5.3	Physical Interpretation	24
1.6	Structure of This Work	27

2	Peak Detection	28
2.1	Assumptions	28
2.2	Synthetic Data Assuming A Top-Hat Response Function	30
2.3	Peak Selection	31
2.4	Estimating the response function	36
3	Peak Shapes	42
3.1	Data Sets	42
3.1.1	Ark 564	42
3.1.2	1H 0707-475	46
3.1.3	NGC 4051	48
3.2	Peak Selection and Measurement Parameters	52
3.3	The Peak Shapes	54
3.4	Peak PSD analysis	55
3.5	The dependence of peak selection on energy band	61
3.6	The effects of non-stationarity on the peak response in Ark 564	63
3.7	The nature of the delayed excess in Ark 564	65
3.8	Explanations for the extracted peak responses	67
3.8.1	Intrinsic flare temporal structure	67

3.8.2	Inverse Compton upscattering time delays	68
3.8.3	Positive lags from propagating fluctuations	68
3.8.4	Energy Dependent Reverberation	70
4	Energy Spectral Analysis of Ark 564	71
4.1	The Peak Energy Spectrum	72
4.2	Lag-Energy Spectra	76
4.3	Modelling the Lag-Energy Spectrum	79
4.4	Physical Interpretation	82
4.4.1	Absorption Model	83
4.4.2	Comptonization Model	87
4.4.3	The Preferred Scenario	91
5	Angular Dependence of Comptonization of X-ray Photons	93
5.1	Compton Scattering	94
5.2	The Monte-Carlo Simulation	97
5.3	Verifying the Simulation Output	100
5.4	Reflection from a hot, infinite slab of finite optical depth	103
5.5	Comparison to Ark 564	112

5.6	Spherical reflectors	116
6	Conclusions and Avenues for Further Research	120
6.1	Further Research	121

1. Introduction

1.1. Active Galactic Nuclei

There is a well established consensus that most, or even all, galaxies contain a supermassive black hole at their centre (Kormendy and Richstone 1995). These objects have masses in the range of $10^6 - 10^9 M_{\odot}$ (Magorrian et al. 1998). Many galaxies contain a highly luminous central engine which is associated with the supermassive black hole (Ferrarese and Ford 2005). This is known as the Active Galactic Nucleus (AGN). The AGN can have a luminosity comparable to, or even outshining, that of its host galaxy, spread over the full range of the electromagnetic spectrum.

AGN can broadly be separated into two categories - the radio-quiet and the radio-loud (Antonucci 1993) based upon the prominence of the radio band emission. These categories are further subdivided based upon the observed properties of the AGN. The main group among radio-quiet galaxies are the Seyfert galaxies, comprising $16 \pm 5\%$ of all galaxies (Maiolino and Rieke 1995). It has been argued that other categories of radio-quiet AGN (Quasars and some Low Ionisation Nuclear Emission-line Regions) are variations on the Seyfert galaxy, with the distinction based on the relative luminosity of the AGN and host galaxy (Antonucci 1993).

Seyfert galaxies are further subdivided into Type-1 and Type-2 galaxies, with intermediate categories such as Types 1.2, 1.5, 1.8, and 1.9. In Type-2 Seyferts the line widths are the same. Type-1 Seyfert galaxies exhibit broad emission lines,

strong featureless continua, and strong X-ray emission in their spectra which are not observed in Type-2 galaxies (Antonucci 1993). The intermediate types exhibit a mix of the features of Type-1 and Type-2 galaxies. Rather than being separate types of objects, the unification scheme of Antonucci (1993) proposes that the observed differences in Seyfert galaxies are due to observational orientation effects. The ratio of Type-2 to Type-1 Seyfert galaxies was found to be 4:1 in a survey by Maiolino and Rieke (1995).

Under the Antonucci (1993) unification scheme, AGN are comprised of three main components. The first is a central engine powered by accretion processes with a length scale of hundreds to tens of thousands of light-seconds. The second is the Broad Line Region, with length scales of light-days to light-weeks. The BLR consists of clouds of gas moving at high velocity in the vicinity of the central Supermassive Black Hole (SMBH) which reprocess the flux from the central engine, giving rise to velocity broadened emission lines. The third component is an optically and geometrically thick torus with length scales on the order of parsecs.

In this scheme, Type-2 galaxies are those which are seen by us in such a way that the central source is obscured by the thick torus. Type-1 Seyfert galaxies are those in which the torus does not obscure view of the central source and BLR. Maiolino and Rieke (1995) used their estimate of the ratio of Seyfert-2 to Seyfert-1 galaxies (4:1) to estimate that 80% of the sky as seen by the central source is covered by the torus, corresponding to an opening angle of $\sim 74^\circ$. They further argue that the intermediate types are due to obscuration of Type-1 AGN by the plane of the host galaxy.

A more recent refinement to this broad scheme is the suggestion by Elvis (2000), in which the torus is replaced by a relatively thin outflow with an opening angle of $\sim 60^\circ$ and a dispersion of 10° . The outflow is composed of material which rises vertically from a small range of radii on the accretion disk, and is then accelerated away from the centre of the system. In this scheme, different classifications of galaxy are observed depending on the orientation of the outflow to the observer.

Narrow-line Type 1 Seyfert galaxies (NLS1) are particularly interesting for X-ray research due to their highly variable nature in the X-ray regime. NLS1 galaxies exhibit features such as a strong featureless continuum and high-ionization emission lines, but lack the broad Hydrogen Balmer lines which typify Type-1 Seyfert galaxies (Osterbrock and Pogge 1985). Boller et al. (1996) note that rather than being a separate category of Type-1 Seyferts, NLS1s are part of a continuum of line widths observed in Type-1 galaxies. Turner et al. (1999) construct a sample by characterising NLS1 galaxies as being those Type 1 Seyfert galaxies in which the FWHM of the Hydrogen- β line is $< 2000 \text{ km}^{-1}$. The narrow spectral lines can be interpreted as being due to two things. The first possibility is that the BLR is proportionally further from the central SMBH. The second possibility is that NLS1 galaxies have lower masses. Boroson (2002) have noted that NLS1 galaxies have low mass accretion rates and high Eddington ratios, indicating that contain lower mass black holes than other AGN.

Boller et al. (1996) note that NLS1s exhibit steeper power-law spectral components, and exhibit rapid variability. Turner et al. (1999) found a significant

correlation between the width of the H- β line and the variability of Type-1 Seyfert galaxies, with NLS1s being an order of magnitude more variable. The variability was calculated by measuring the excess variance method of Nandra et al. (1997). It is this rapid X-ray variability which makes NLS1s so important for the study of the inner regions of AGN. Light travel time arguments state that (in the absence of relativistic effects), an object cannot vary on timescales shorter than its light-crossing time (Terrell 1964; Rees 1966). The rapid variations observed indicate that the X-rays observed come from the very innermost regions of the AGN, and so these objects are extremely important for the study of accretion processes.

1.2. The Components of the Central Engine of Seyfert Galaxies

1.2.1. *The Black Hole*

The SMBH at the centre of the AGN can have a mass of between 10^6 and 10^9 solar masses, with the largest outlier being a blazar with an estimated mass of $4 \times 10^{10} M_{\odot}$ (S5 0014+813, Ghisellini et al. (2009)). The mass of the SMBH gives rise to a useful length scale, the Gravitational Radius, defined as $R_g = \frac{GM}{c^2}$. A useful rule-of-thumb is that for a $10^7 M_{\odot}$ black hole, $R_g \approx 50$ light-seconds. For a non-rotating black hole, the event horizon is at the Schwarzschild Radius $R_s = 2R_g$. For a maximally rotating Kerr black hole it is at R_g .

1.2.2. *The Accretion Disk*

The accretion disk is the structure which arises due to the need for accreting matter to shed angular momentum. Without such a structure, infalling matter would not enter the black hole but instead continuously orbit. The accretion disk transports angular momentum through viscous effects, heating the disk and removing energy through radiation. The standard model for such an accretion disk is the optically thick, geometrically thin disk of Shakura and Syunyaev (1973). Matter slowly moves inwards in steadily decreasing orbits until it reaches the Innermost Stable Circular Orbit (ISCO) of the SMBH, at which point it rapidly falls into the black hole.

For a non-rotating black hole the ISCO is at $6R_g$, while for a rotating Kerr black hole the ISCO can be as low as $1.2R_g$ for prograde orbits (Abramowicz et al. 2010; Fabian et al. 2014). The technique of Fabian et al. (2014) models the gravitational redshift of the Fe-K emission line profile, and by assuming that the Fe-K emission is mostly from the ISCO determines its location. This is then used to estimate the spin of the black hole.

The thin-disk model has the temperature of the disk greatest at the smallest radii. An interesting consequence of this is that smaller black holes, with smaller ISCOs, have the peak in radiation from the disk at higher temperatures. This means that while the black-body disk radiation from stellar mass black holes peaks in the X-ray regime, the black-body component of radiation from AGN is in the ultraviolet, a notoriously difficult area to study.

1.2.3. *The ‘Corona’*

The so called ‘corona’ is the most mysterious aspect of AGN. A common feature in X-ray spectra is the power-law spectrum, and the mechanism by which the power-law is generated is assumed to be Compton upscattering of soft seed photons by a region of hot gas. Upscattering by hot gas can easily generate power-laws from arbitrary seed spectra (Sunyaev and Titarchuk 1980). The population of hot electrons required for such upscattering is ascribed to a region known as the ‘corona’. This term is used to draw an analogy to the extremely high temperatures of stellar coronae compared to their surface, and does not propose that the same mechanism is at play. What is unknown is the geometry and location of this upscattering region. Haardt and Maraschi (1991) suggest a Comptonizing ‘coronal’ layer sandwiching the disk (see also Sunyaev and Titarchuk (1980)). Done et al. (2012) suggest a model in which the outer regions of the accretion disk radiate as a black-body, and the inner regions as a black-body with a small amount of upscattering. The inner regions are further surrounded by a hot skin which provides the upscattering to produce a high-energy power-law. This forms the basis of the XSPEC model `optxagn`. Section 5 of this work proposes a model in which the Comptonizing region is a clumpy distribution extending 10s to 100s of R_g from the centre of the system.

1.2.4. *Absorbers*

Absorption is prevalent enough in X-ray astronomy to warrant the existence of the XSTAR program (Kallman 2005), which calculates the absorption and emission spectra which arise from the passage of arbitrary input spectra (defined at the least between 0.0136 keV and 13.6 keV) through uniform, circularly symmetric slabs of gas. The topic of X-ray absorption has been given a lengthy and comprehensive treatment by Turner and Miller (2009). What is of relevance here is the location of these absorbers. The proposal of Miller et al. (2007) is that the variability in the absorption exhibited by Mrk 766 is due to partial covering of the X-ray source by a clumpy, extended, distribution of absorbing gas, with the covering fraction varying in time. The origin and nature of such an extended distribution is uncertain, and has plenty of scope for interesting ideas. Concepts include clumpy winds and outflows as well as infalling matter.

1.2.5. *Wind*

Winds of bulk outflowing material are often observed in association with AGN, and can be important on a wide range of length scales. As well as X-ray evidence (Tombesi et al. 2013), Broad Absorption Lines in the UV provide further evidence that outflowing material is common in AGN (Proga 2007). Tombesi et al. (2010) Found evidence for relativistic outflows in 35% of Type-1 Seyfert galaxies surveyed, and Tombesi et al. (2013) reported that $> 60\%$ of Type-1 Seyferts have large-scale warm absorbers - parsec to kiloparsec scale winds. On kiloparsec scales, winds can

be a factor in galactic evolution (Faucher-Giguère and Quataert 2012). On shorter length scales of $\approx 100R_g$ outflows can be an important aspect of the behaviour of the innermost regions of AGN, and Tombesi et al. (2013) claim that the ultra-fast outflows and large-scale warm absorbers are part of the same unified structure. Proga et al. (2000) have performed hydrodynamical simulations of the launch of radiatively driven winds, demonstrating that their structure can be very complex. Leighly and Moore (2004) and Leighly (2004) have modelled the UV spectrum of 1H 0707–495 as being due to a wind with outflowing velocity 2500 km s^{-1} . Gupta et al. (2013) claim the existence of a relativistic outflow in Ark 564 with velocities of $\approx 0.1c$.

1.3. Probing the Central Regions

Going beyond the simplistic unification model is a difficult task. The small size of regions involved, and the great distance of their host galaxies from Earth, mean that the Broad Line Region and the central source are not resolvable by current telescopes, and are unlikely to be so within any reasonable timeframe. Estimating the geometrical properties of the central regions requires more sophisticated methods.

A technique which has seen significant success is optical reverberation mapping, which can be used to find the size of the BLR and to deduce the mass of the SMBH (Peterson 1993). Reverberation mapping uses the delay between variations in the continuum flux and variations in the BLR emission to estimate the travel time from the source of the continuum radiation to the BLR clouds. The velocity of the

BLR clouds is known from the broadening of the emission lines. Combining the two pieces of information allows an estimate of the SMBH mass to be reached.

To investigate regions even closer to the central SMBH, the observations in the X-ray regime are required. Light travel time arguments hold that (in the absence of relativistic effects) an object cannot vary on timescales shorter than the light-crossing time. The X-ray flux varies on all timescales from hundreds to millions of seconds. This means that X-ray observations can be used to explore the very inner regions of the AGN, potentially down to the inner edge of the accretion disk (Fabian et al. 2012).

The techniques used with X-ray observations fall into two main categories - those which examine the temporal behaviour, and those which examine the spectral behaviour. The observations produce a list of photons captured by the X-ray telescope, and it is up to the investigator to choose how to slice the data. A timing analysis groups photons into broad energy bins to achieve good statistics with high time resolution. Conversely, a spectral analysis averages over long periods of time to obtain a high energy resolution. With long observations of particularly luminous objects it is possible to pursue a hybrid approach, accounting for both spectral and temporal properties of the data, albeit both at lower resolution than an independent analysis.

1.3.1. *Timing Analysis*

The techniques used for timing analyses of AGN in the X-ray make considerable use of the Fourier transform of the time series. This has the advantage that the observed time series can be viewed as a combination of variations on different timescales. Noting the frequency at which features appear in the Fourier transform allows estimates of physical properties to be reached. The two main techniques used are estimates of the Power Spectral Density (PSD, or Power Spectrum) and the Lag Spectrum. The Power Spectrum looks at the amplitude of variations on different timescales, while the Lag Spectrum looks at the phase component of the Cross Spectrum between two time series. The Power Spectrum is a widely used tool across many fields and a brief explanation of its application to X-ray observations of AGN is given in Section 1.4. The Lag Spectrum is more of a peculiarity and is explained further in Section 1.5.

There are difficulties in using such techniques, however. The nature of the observations means that the time series obtained are noisy and full of gaps. Noise is ever present in practical observations, and all the more so when using data from objects with low count rates. The gaps in the time series potentially come from several sources. On the longest timescales, the dataset can be formed from separate observations with separations of days or weeks. On medium timescales, there can be gaps due to the orbital period of the telescope used. On the shortest timescales are unpredictable gaps due to excising periods of high flaring background flux.

These inevitable flaws in the data make standard methods of calculating the Fourier transform useless - the Fast Fourier Transform (FFT) relies on having a

complete timeseries, as do other exact methods. Fortunately, inexact methods can provide useful information, and with care can be used on even the most bizarre time series. The Lomb-Scargle Periodogram was developed for just this problem, which is common in astrophysical datasets (Lomb 1976; Scargle 1982), and provides an estimate of the Power Spectrum. Scargle (1989) claims to extend this to cross-correlation between unevenly sampled time series, though this method is little known and has met with significant difficulties in practical implementation. The downside of the Lomb-Scargle method is that it does not provide confidence regions.

This downside is corrected by the Maximum-Likelihood method developed by Miller et al. (2010b) and tested and verified by Zoghbi et al. (2013). This method simultaneously fits to the maximum-likelihood values of the power and cross spectra, then steps around the optimum values to give 68% confidence regions for each frequency bin. The likelihood is calculated by creating a model covariance matrix C from the proposed power and cross spectrum values in each frequency bin. Minimising $(x - \bar{x})C^{-1}(x - \bar{x})^T$, where $(x - \bar{x})$ is the vector of mean-subtracted observational data points, maximises the likelihood. Some care should be taken with the method. It can be computationally intensive, so the size of the frequency bins should be chosen with care by considering the computation time, the quality of the data and the desired resolution of the output.

1.3.2. *Energy Spectral Analysis*

Analysis of the energy spectrum principally involves comparing the data to model spectra generated by a tool such as XSPEC (Arnaud 1996). This program, when used with the correct detector response files, is capable of modelling the photon counts in each of the instrument detector channels. These model counts can then be compared to the observed counts to provide and display the model and observed spectra.

A complicating aspect of this process is that it is not possible to extract the ‘true’ spectrum from the instrument data and response alone. The instrumental response has a convoluting as well as a multiplicative effect on the incoming photons. This can easily be modelled in the forward direction, estimating the output of the detector for a known input spectrum, but the reverse is not true. To get a good estimate of the true spectrum from the object observed requires a model that fits as closely as possible. The effect of estimating the ‘true’ spectrum with a poor model is especially large in the region of sharp features in the spectrum, where convoluting effects would be greatest.

The XSPEC program provides a great many model components which can be used to provide a fit to the data. A number of the model components are different ways of achieving and approximating the same effect. For instance, a correction for absorption by the interstellar medium along the line of sight can be provided by the models `tbabs` and `wabs`, using photoelectric cross-sections by Wilms et al. (2000) and Morrison and McCammon (1983) respectively. Other model components include power-laws, black-body spectra, line emission, and absorption.

These model components are combined in ways that make physical sense to reach conclusions about possible explanations for the observed spectra. For instance the researcher might choose a model for the continuum power-law emission such as `optxagn` or `nthcomp`, add in a component to represent the black-body disk emission (e.g. `bbody`), and multiply the ensemble by an absorbing component (generated with the tool `XSTAR`).

1.4. Power Spectra

The PSD plots the amplitude-squared of the Fourier transform against frequency, and for a stationary, gaussian distributed time series contains the full statistical information. A time series $s(t)$ has a Fourier transform $S(\nu)$. The power spectrum is found by

$$P_S(\nu) = |S(\nu)|^2 = S(\nu)^* S(\nu) \quad (1)$$

It is common to observe that the PSD of Seyfert galaxies appears to take the form of a broken or bending power-law (McHardy 2010), with the power-law index below the break frequency taking a value of $\Gamma \approx -1$. For this reason, it is common to plot $\nu P(\nu)$ against ν . However, there are alternative approaches to fitting simple functions to PSDs. McHardy et al. (2007), for instance, fit two broad Lorentzian profiles to the power spectrum of Ark 564.

There are a few aspects to be aware of when estimating the power spectrum. Firstly, the finite binning in the time domain is effectively a convolution of the ‘true’

time series with a top-hat function of width equal to the bin width. Therefore the power spectrum of the binned time series must be divided by the power spectrum corresponding to the top-hat. Secondly, the presence of instrumental counting noise in the data set must be accounted for. This can be corrected by simply subtracting a white noise power spectrum of the appropriate amplitude.

Thirdly, the time series observed is effectively a multiplication of a lengthy, continuous ‘true’ time series with a window function in the time domain (Harris 1978). This window function takes a value of zero at all times except for those in which data is present. This could be as simple as a single top-hat function with a width corresponding to the observation length, or a complicated function encompassing multiple observations, periodic dead-time due to the telescope orbit, and individual excised time bins due to high flaring background flux. The effect of the multiplication in the time domain is a convolution in the frequency domain. This means that the observed power spectrum is a convolution with the spectrum corresponding to the window function. Typically the window function is very wide, resulting in a narrow function in the frequency domain. This means that only the very lowest frequency bins will be affected when logarithmic binning is employed. The negative power-law nature of many power spectra leads this phenomenon to be sometimes known as ‘red noise leak’, and the effect is to shift power to higher frequencies.

1.5. Lag Spectra

1.5.1. Definition

The *Lag Spectrum* is used to visualise the phase information between two time series in Fourier space (Nowak et al. 1999). In X-ray astronomy these time series are usually two broad energy bands denoted as the soft (low energy) and hard (higher energy) bands. The time series $s(t)$ and $h(t)$ have Fourier transforms $S(\nu)$ and $H(\nu)$. These Fourier transforms give rise to the Cross Power Spectrum

$$C(\nu) = S^*(\nu)H(\nu) \quad (2)$$

The amplitude of $C(\nu)$ contains the information required to calculate the coherence (Nowak et al. 1999). The phase $\Phi(\nu) = \text{Arg}[C(\nu)]$ conveys additional information about the temporal relationship between the two bands. The phase is further interpreted as $\Phi(\nu) = 2\pi\nu\tau(\nu)$.

The lag spectrum is a plot of $\tau(\nu)$ against ν with the convention that a positive value of τ indicates that the hard band component with frequency ν lags the soft band component by τ seconds. Fig. 1 shows the lag spectrum for 1H 0707–495 between the soft (0.3-1.0 keV) and hard (4.0-7.5 keV) bands (reproduced with permission from Miller et al. (2010a)). The lag spectrum was estimated through the maximum-likelihood method of Miller et al. (2010b). Also shown is a fit using a simplistic top-hat reverberation model (see Section 1.5.3).




Fig. 1.— The lag spectrum of the narrow-line Seyfert I galaxy 1H 0707–495, reproduced with permission from Miller et al. (2010a). The soft band consists of photons with energies between 0.3 keV and 1.0 keV, while the hard band is 4.0 – 7.5 keV. Bars give the frequency range of each bin and the 68% confidence intervals of the lags. The solid line is a fit with a simple top-hat reverberation model.

1.5.2. History

The lag spectrum was first used by van der Klis et al. (1987) to examine the low mass X-ray binaries Cyg X-2 and GX 5-1. Previous studies made use of the Cross-Correlation Function (CCF) to examine the relationship between energy bands. The move to Fourier analysis was motivated by a desire to examine the lag in a specific frequency range corresponding to the Quasi-Periodic Oscillation

frequently observed in X-ray binaries. The technique quickly became a standard tool in the analysis of X-ray binaries, with both positive and negative lags routinely observed (Miyamoto et al. 1988; van der Klis et al. 1987). It should be noted though that the negative lags are rarely significantly so, and/or occur at frequencies where there is reason to suspect other effects. For instance the van der Klis et al. (1987) lag spectra show negative lags at the lowest frequencies. This is the region which can be suspected of being affected by red-noise leak.

The first use in Active Galactic Nuclei came later with an observation of NGC 7469 by Papadakis et al. (2001). The delay in the extension of the technique to AGN was due to the poor quality of the data available. The NGC 7469 analysis made use of > 2.76 Ms of *RXTE* observation time. With newer telescopes such as *XMM-Newton*, acceptable lag-spectra can be formed from as little as 40 ks of data, though longer observations are preferable.

The first unambiguous observation of negative lags in AGN was in the object 1H 0707-495 (Fabian et al. 2009; Zoghbi et al. 2010), though there had been tentative observations made previously. The original NGC 7469 lag spectrum had an (insignificant) negative lag in a single frequency bin (Papadakis et al. 2001). A CCF analysis of IRAS 13224-3809 had identified negative lags, albeit without the frequency information (Gallo et al. 2004). McHardy et al. (2007) found hints of negative lags at high frequencies in Ark 564, since confirmed by Legg et al. (2012).

There have since been numerous detections of negative lags. De Marco et al. (2013) formed lag spectra for all 32 AGN in the CAIXAvar catalogue which had published mass estimates, observation lengths greater than 40 ks, and exhibited

variability. Of these 32 AGN, 15 exhibit negative (‘soft’) lags. These 15 AGN fall into two families - eight in which the the low frequency lags are positive, and six in which they are negative, with one ambiguous. In each case the lag spectrum has a characteristic frequency above which it switches to the opposite of the low frequency behaviour. A lengthy observation of NGC 4051 has been claimed to exhibit switching between the two behaviours within a single object (Alston et al. 2013).

1.5.3. Physical Interpretation

The physical interpretation of lag spectra in AGN has been a controversial topic. The information they provide gives a constraint to any models that seek to explain the behaviour of AGN in the X-ray regime. While any model must be able to explain the observed lag spectra, there are multiple models which have been developed to do so. Unfortunately the lag spectra in themselves are not enough to distinguish between the competing models.

Perhaps the most popular model is the Propagating Fluctuations model (Lyubarskii 1997; Arévalo and Uttley 2006; Zoghbi et al. 2010). In this model, density fluctuations propagate inwards on the accretion disk. The inner regions of the disk produce a harder source spectrum. An increase in activity produces a soft spectrum which hardens with increasing time as the fluctuation moves to the inner regions of the disk. The disk at a particular radius is assumed to produced fluctuations on a timescale characterised by the local viscous frequency. Longer timescale variations are produced at larger radii. Since this model alone is

incapable of producing negative lags, additional soft band reverberation from the inner edge of the accretion disk is invoked to explain negative features in the lag spectrum. The model has seen a great deal of use in the study of X-ray binaries.

The propagating fluctuations model runs into difficulties in explaining X-ray observations of AGN. While the accretion disk of an X-ray binary radiates directly in the X-ray, the larger mass of the SMBH in AGN means that the disk radiation is in the UV. This in itself is not fatal to the model - one can imagine a ‘corona’-like skin over the disk to upscatter photons close to their point of origin. Such a scheme would, however, be unable to produce negative lags. Instead, the light-bending model of Miniutti and Fabian (2004) is invoked. In this model, a hotspot situated at a variable height along the axis of the system illuminates the disk, with gravitational effects increasing the influence of reflection from the inner edge of the disk, which gives rise to negative lags. It is difficult, however, to reconcile this with the propagating fluctuations. The hotspot cannot be the upscatterer of the UV photons, since Compton upscattering erases the ‘memory’ of the input spectrum (see Section 5). In Miniutti and Fabian (2004) the variability is entirely due to variations in the height of the hotspot above the disk plane, with no influence from propagating fluctuations. Zoghbi et al. (2010), however, invoke the hotspot model alongside propagating fluctuations.

The main competing model views the lag spectrum as primarily due to the shape of the hard band response function (Miller et al. 2010a). Sharp, asymmetric features in the response function produce negative features in the lag spectrum when Fourier transformed. Miller et al. (2010a) demonstrated that simplistic

top-hat response functions could easily produce features reminiscent of the negative lags observed 1H 0707–495. The proposal of the model is that the shape of the lag spectrum is explained by reverberation of the source spectrum by clouds of surrounding gas.

1.6. Structure of This Work

In Section 2, a novel method is developed for detecting and analysing peaks within the gappy time series associated with X-ray observations of Active Galactic Nuclei.

In Section 3, three observations of AGN are introduced and the new peak detection method is applied to them. The peak shapes are estimated and their properties compared to the properties of the overall time series. Particular attention is paid to the object Ark 564 due to the striking results of the peak analysis, which show a peak and delayed hard band excess characteristic of reverberation.

In Section 4, the object ARK 564 is further analysed in an attempt to simultaneously explain both the temporal and energy spectral data. Energy spectra corresponding to the peak and delayed excess are estimated and used to put physical limits on reverberation models. The preferred model is one in which the delayed excess is due to a reverberation by a hot comptonizing region.

Section 5 addresses the shortcomings of current simulations and models of hot Comptonization. Research to date has not presented the outcome of reflection of X-rays by a hot Comptonizing region. A Monte-Carlo simulation is used to rectify this, and the full angular dependence of output radiation is shown for a range of geometries.

2. Peak Detection

2.1. Assumptions

The task of detecting peaks within the X-ray time series is undertaken with a set of assumptions in mind. The first of these is that the time series in each energy band consists, at least in part, of a common time series convolved with a response function. This response function has the potential to be energy-dependent. The second is that the time series contain short timescale flares which might enable the identification of the signature of the reverberation in each energy band.

The justification for first of these assumptions is mathematical. Any time series can, in principle, be decomposed into an underlying time series of source variations, convolved with a response function

$$F^{obs}(t) = F^{src}(t) * R(\tau) \tag{3}$$

where F^{obs} is the observed flux as a function of time t , F^{src} is the time flux at the source of the variations, and $R(\tau)$ is the response as a function of delay τ . The response function has occasionally been referred to in the literature as the ‘transfer function’ (Nowak et al. 1999). In this document I shall stick to the convention that the ‘response function’ refers to the time domain, while the ‘transfer function’ refers to its Fourier transform when the convolution is considered in the Fourier domain.

$$f^{obs}(\nu) = f^{src}(\nu)T(\nu) \tag{4}$$

From the Fourier transform (in terms of frequency ν) we can see that any time series can be decomposed into a source time series and response function, and that without further information or constraints the decomposition is degenerate. The degeneracy, and guaranteed existence, of this decomposition means that it must be examined closely in light of physical models.

The second assumption is that the time series contain features which allow the estimation of the response functions. The ideal, albeit unphysical, scenario would be one in which the underlying time series contained a series of isolated δ functions. This would allow the easy extraction of the response functions. A more plausible scenario is that the underlying variations contain sharp peaks which come close to the ideal case. This is justified by a qualitative examination of the X-ray time series, which do indeed contain sharp peaks. The assumption is required due to practical considerations making all available methods of blind deconvolution unfeasible.

Blind deconvolution is the task of obtaining the two convolved functions given only their convolution and a set of domain-specific constraints on the nature of the convolving functions. This difficult task has been given a lot of attention in the field of image processing, and good overviews of available techniques are given by Kundur and Hatzinakos (1996) and Takahata et al. (2012). However, the practical nature of the problem at hand, namely the presence of significant amounts of noise in the time series, immediately rules out the easiest available methods. The fact

that the time series contain gaps means that obtaining a Fourier transform is non-trivial, and hugely complicates the remaining methods. In principle it would be possible to fit to a model light curve F^{src} and model response functions, but the vast number of parameters makes this computationally impractical. The method that remains is to identify peaks in the time series, and make the assumption that the underlying source behind those peaks approximates a δ function added to some baseline curve.

2.2. Synthetic Data Assuming A Top-Hat Response Function

To aid in testing and verifying the peak detection algorithm, a synthetic dataset is required. This should have known properties and mimic the observational data. By running the algorithm against the synthetic data it is possible to ensure, first, that known features are recovered and, second, to identify any biases and inaccuracies in the algorithm’s output.

The synthetic data is based on a time series $F_{synth}^{src}(t)$, which is generated from a broken power-law PSD with a break frequency of $\nu_{break} = 7 \times 10^{-4}$ Hz. The power-law index at frequencies below the break frequency is $\alpha = -1.1$, while above the break it is $\beta = -3.0$. The synthetic time series consists of 3500 time bins each of width 128 s. The synthetic time series at this point is a Gaussian process, with a normal distribution, whereas an examination of observational X-ray time series shows them to have lognormal distributions (Section 3.1). To more closely approximate observational data, the synthetic data is exponentiated ($F'(t) = \exp[F(t)] / \exp[1.0]$).

A second time series is then generated by assuming convolution with a top-hat response function $R_{\square}(\tau)$, mimicking a reverberation scenario. The centre of the top-hat is 1920 s after the δ function and the width is 1280 s. The top-hat integrates to a fraction $f = 1.0$ of the integral of the δ function. The convolved series is generated as

$$F_{synth}^{reverb}(t) = F_{synth}^{src} * R_{\square} \quad (5)$$

These time series are, at this stage, normalised with a mean of unity. To generate useful synthetic data this mean was multiplied by 3000 to give an average count rate of 23.4 s^{-1} , and Poisson noise applied. This count rate is in line with that from a luminous AGN - for instance the 0.4-2.0 keV flux from Ark 564 is 24.0 s^{-1} .

2.3. Peak Selection

The task is to identify peaks in the time series to use in the analysis. The most naïve method for achieving this would be to simply examine each point in the time series, and test whether it is greater than its neighbouring points. It is clear that this would not be a satisfactory solution, since the presence of noise in the observed time series means that there are a large number of local maxima that amount to little more than statistical variations. There is also the factor that the time series itself is formed from a binned list of delta functions (the photon arrival times), and different choices of window length would give a very different set of peaks. Instead, we wish to identify times which correspond to peaks of some significance above a

local baseline flux.

The local baseline flux level consists of variations on timescales much longer than the characteristic width of the short timescale peaks. One method to extract this might be to use a high-pass filter in the Fourier domain, but the presence of gaps in the observational time series makes this difficult to implement. While the Lomb-Scargle periodogram is capable of estimating the PSD of a time series with gaps, it loses the phase information required for a high-pass filter. The method chosen here is to use a median filter.

A median filter of width N returns a value at each point in the time series equal to the median of the N time bins centred on that point. When the filter window overlaps a gap in the time series the missing value can simply be ignored. The result is two time series - a median filtered baseline, and a time series of residuals. The median filter has the property that it preserves features with a length scale of less than half the filter window.

For instance, consider a time series which is uniformly zero with a superimposed set of isolated features with width no greater than 10 time bins. A median filter with $N = 21$ will result in a baseline which is, as expected, zero at all points, and the shape of the features preserved in the residuals. This is in contrast to a Gaussian smoothing function, which will only preserve features if their values average to the baseline value. Since the observational data has a non-Gaussian distribution this is not the case, and a median filter is preferred.

Having subtracted the median filtered baseline, the residual time series is inspected to identify peaks. Since the observational data contains Poisson




Fig. 2.— (Left) The PSD of the synthetic time series, plotted as $\nu P(\nu)$. (Right) The expected peak shape of the unmodified source time series, normalised to a peak height of unity.

distributed shot noise, it is appropriate to smooth the residual time series. The choice of smoothing function is an important one. A smoothing function will tend to emphasise features with a shape similar to that of the smoothing function. Since we wish to identify peaks in the time series, we can make use of the information within the time series to find the optimal smoothing function.

This optimal smoothing function is obtained from the PSD of the time series by assuming symmetry and a single δ function at $t = 0$. A smoothing function which is wider than the ‘true’ underlying variations would confuse and blur signals in the time series. A function which is too narrow would over-emphasise noise.

The optimum solution is to use the characteristic length scale of the dataset, taken from the PSD. The PSD and associated shape are shown for the synthetic data in Fig. 2.

Having smoothed the series, peaks are found by identifying points which are greater than both of their nearest neighbours, and have a statistical significance of greater than 2σ . The statistical significance is found through propagating the shot-noise errors through the smoothing process, and maxima selected as peaks if their value is greater than 2σ .

The steps of this process are shown in Fig. 3 for a section of the synthetic light curve, and are summarized as follows:

- A median filter of width $N = 21$ is used to identify a baseline time series.
- The baseline is subtracted from the raw time series to give a residual time series
- The residual time series is smoothed with a smoothing function derived from the power spectrum (Fig. 2, right).
- Times corresponding to maxima with a value more than 2σ greater than zero are selected as peaks.

The choice of the median filter width is based on the bin size and the width of the smoothing function, and is chosen to be odd for the practical purpose of simplifying the process of finding the median value. The significance criterion, set at 2σ , is chosen to provide sufficient peaks worthy of analysis. Setting this




Fig. 3.— (Top) A section of the raw synthetic light curve in black, and the median filtered baseline in grey. The light curve has been normalised to the series mean. (Middle) The residual light curve after subtraction of the median filtered baseline. (Bottom) The smoothed residual curve, with detected peaks marked with vertical lines. For clarity, the synthesised errors are not shown.

parameter too high will result in a lack of selected peaks, especially in the case of poor quality noisy data. Setting the parameter too low will select peaks that are not isolated enough from the baseline for the analysis to work.

2.4. Estimating the response function

Having selected a set of times associated with peaks, we now seek to examine the shape of those peaks at longer timescales to gain insight into the nature of the response function. The low frequency regions of the lag spectra can exhibit lags of order 1000s, which we wish to capture. In that case, the current baseline subtraction is inadequate, since it removes the features we are looking for. The solution is to identify a baseline with a longer filter length.

One possibility is to use a simple median filter with a larger window length (e.g. $N = 81$). However, a visual examination of the resulting baseline light-curves reveals that they exhibit features which are likely inappropriate. An example of a single-stage baseline is shown in the bottom panel of Fig. 4, with the sort of inappropriate feature produced by the single step filter seen at ≈ 140 ks. Short timescale features in the subtracted baseline will have an undesirable impact on the residual time series. The alternative is to use a multistage filtering process. In the multistage filtering process the raw light curve is first filtered with $N = 21$ to produce a first stage baseline. This baseline is the same as that used for the peak detection. A second stage baseline is produced by filtering the first stage with $N = 41$, and the third stage by filtering the second with $N = 81$. The three stages are shown in the upper panel of Fig. 4.

Having generated a baseline light-curve, this is subtracted from the raw light-curve to form a time series of third-stage residuals. It is these residuals which are used to form an estimate of the response functions. Each of the peaks found with the algorithm of Section 2.3 has a shape given by $N = 81$ time bins from the




Fig. 4.— (Top) A section of the raw synthetic light curve in black, and the first ($N = 21$), second ($N = 41$) and third ($N = 81$) stage median-filtered baselines in red, green and blue respectively. The light curve has been normalised to the series mean. (Bottom) The raw light curve alongside the third stage baseline (blue) and a single stage median filter with $N = 81$ (red). Note the inappropriate feature in the single stage baseline at 140 ks.

residual time series centred on the peak time.

A key feature of this method is that the peak times are identified in a single band (the soft 0.4 – 1.0 keV band) and are used to generate response functions in all other bands. Because we are interested in the relative shapes of the response functions, we scale by height at $t = 0$. This characteristic height is found by

smoothing the third-stage residual time series and taking the value at the peak time. The smoothed residual light curve is used to reduce the impact of shot noise on the central peak value, while having a consistent measure of the peak height.

To estimate the response functions the weighted average of the peaks is taken, with weights given by the inverse of the total shot-noise variance of the individual peaks in the soft band. This averaged shape is scaled such that the amplitude at $t = 0$ is unity for ease of visual comparison. The use of the soft band for peak detection and weighting ensures that the estimated response function in the other energy bands is completely unbiased.

The algorithm is tested against the synthetic time series generated in Section 2.2. This allows verification that the algorithm works, and helps identify biases and weaknesses. The results of this test are shown in Fig. 5. The response functions applied to the synthetic series are shown in black. The theoretical peak shape expected from the PSD after convolution with the response function is shown in blue. The shapes estimated by the algorithm are shown in red. Also shown, in green, is the result of a convolution of the soft band estimated shape with the hard band response function. This shows that the algorithm convincingly separates the long timescale baseline and reproduces the signature of the top-hat response function. All shapes have been normalised to have a value of unity at $t = 0$.

The most obvious difference between the theoretical peak shapes (blue) and those given by the algorithm (red) is that the theoretical shapes are non-zero on long timescales. To an extent this is expected, as the baseline subtraction removes long timescale features. After this effect has been accounted for, the estimated hard




Fig. 5.— Application of the algorithm to the synthesised time series data. (Left) The soft band response, a δ function, in black with the theoretical peak shape in blue. The peak shape given by the algorithm is in red. (Right) The hard band response function in black with the shape estimated by the algorithm in red. The convolution of the theoretical soft band peak shape with the hard band response function is in blue, while the convolution of the measured soft band peak is in green.

band shape is a good match for the result of a convolution of the soft band shape with the response function. The only issue is a slight offset, with the estimated hard band shape tending to lower values. Hereafter, the extracted shapes shall be referred to as the ‘peak response’, since they show the response of the time series to a peak in the soft band.

The steps of the algorithm for estimating the peak response function can be summarized as follows:

- Identify a set of peaks in the time series with the algorithm of Section 2.3.
- Filter the raw time series in each energy band with a median filter of window length $N = 21$ to produce a first stage baseline.
- Filter the first stage baseline with filter length $N = 41$ to produce the second stage baseline.
- Filter the second stage baseline with filter length $N = 81$ to produce the third stage baseline.
- Subtract the third stage baseline from the raw time series to produce a residual time series.
- For each peak in the set of identified peaks, use a section of the residual time series $N = 81$ bins in width as the shape of that peak.
- Smooth the soft band third stage residuals with the same smoothing function as in the peak detection algorithm.
- For each peak, take the value of the smoothed residuals at the peak time as the characteristic height.
- Divide the peaks by their characteristic height.

- Perform a weighted average of the individual peak shapes, using the total shot-noise variance of the peaks as weights to obtain the estimated peak response.
- Use the standard deviation of the values of the individual peaks in each time bin as the estimated error.
- Scale the averaged shape and errors such that the value at $t = 0$ is unity.

The algorithm is all but guaranteed to identify peaks in the soft band if the level of noise is low enough. The use of a symmetric smoothing function in peak detection and normalisation does not bias the shape of the soft band peak responses, but could have effects on the normalisation if the ‘true’ peak response is asymmetric. If the ‘true’ soft band response is asymmetric, then the peak selection algorithm will select a peak time which is biased away from the ‘true’ peak time. This bias would be consistent across all peaks, and would manifest itself as an offset of the estimate peak response from the nominal $t = 0$. The shape itself would be unaffected. The shape in the hard band is completely unbiased by the soft band results other than this potential offset. Any apparent features are either genuine or the result of statistical variations.

3. Peak Shapes

3.1. Data Sets

Datasets from three Seyfert I galaxies (Ark 564, 1H 0707–495, and NGC 4051) are used in this work. All three galaxies are analysed in Section 3, while Ark 564 is examined in further detail in Section 4.

3.1.1. *Ark 564*

The narrow-line Seyfert I galaxy Ark 564 (redshift $z = 0.0247$, Hewitt and Burbidge (1991)) contains a supermassive black hole with a mass estimated at $2.61 \pm 0.26 \times 10^6 M_{\odot}$ by Botte et al. (2004). This estimate was formed using the Kaspi et al. (2000) relationship between the radius of the Broad Line Region and the luminosity of the AGN at 5100 Å. McHardy et al. (2007) modelled the power spectrum in terms of two Lorentzian components, interpreting the associated lags as supporting the propagating fluctuations model.

A high quality observation of Ark 564 was performed in 2011, and first presented by Legg et al. (2012)¹, with further spectral analysis presented by Giustini et al. (2015). The data set consists of eight separate observations by *XMM Newton* with a total of 486 ks of observation time. The data used here was collected by the EPIC pn detector in Small Window mode. The data was processed

¹The data reduction was carried out principally by T.J. Turner of UMBC, and described here for completeness

with HEASOFT v. 6.12 and SAS v. 11.0.0. Source counts were obtained from a circular region of radius $36''$ centred on the object, filtered to retain only the best quality events (Singles and Doubles², $\text{PATTERN} \leq 4$). Background regions were two boxes $61'' \times 100''$ and $51'' \times 51''$ respectively, filtered to retain only Singles in the 10–12 keV band.

Time intervals in which the background was strongly flaring were removed using a threshold of 0.1 count s^{-1} , leaving 286 ks of ‘good’ time. The source events were then analysed to check for pile-up using the SAS task *epatplot*. Pile-up occurs when there is a high flux of photons hitting the detector. If two photons hit the same pixel in quick succession they might be registered by the detector as an event with the combined energy of the photons. If they hit adjacent pixels they would be registered as a Double (or higher order) event instead of two Singles. This can be checked for by comparing the observed ratio of Doubles and Singles to the theoretical ratio, which is achieved using *epatplot*. The deviations observed of 3-5% were judged to be too high. To reduce the impact of the pile-up, the source region for the four observations with the highest flux was changed to an annulus, excising the central regions of the Point Spread Function where the flux was highest.

For the purposes of time series analysis, the events are binned temporally in bins 128 s in width beginning with the first ‘good’ time of the observations. The

²Singles are those events where the above-threshold charge on the detector is concentrated in a single pixel, and are assigned $\text{PATTERN}=0$. Doubles are where the pixel with peak charge has a single adjacent pixel with above-threshold charge, and are assigned $\text{PATTERN}=1-4$ depending on the direction of the adjacent pixel.

counts in each bin are corrected for any excision performed and the amount of ‘good’ time within the bin period, and bins with less than 50% ‘good’ time are discarded.

The photons are further binned into four broad energy bands - Band 1 (0.4-1.0 keV), Band 2 (1.0-2.0 keV), Band 3 (2.0-4.0 keV) and Band 4 (4.0-7.5 keV). The result is four gappy time series suitable for use with the maximum-likelihood code of Miller et al. (2010b). The gaps result from the long intervals between the eight observations, and the time bins excluded due to periods of high flaring background flux.




Fig. 6.— Histograms of $F_{band} / \langle F_{band} \rangle$ for each of the four broad energy bands of Ark 564.

The mean fluxes in each of the four bands are 18.0 s^{-1} , 6.0 s^{-1} , 1.3 s^{-1}




Fig. 7.— (Right) Power spectra of Ark 564 in Band 1 (Black) and Band 4 (Grey), plotted as $\nu P(\nu)$. (Left) Lag spectrum between Band 1 and Band 4, positive values indicating that Band 4 lags behind Band 1.

and 0.42 s^{-1} for Bands 1, 2, 3, and 4 respectively. For convenience and ease of comparison between bands the fluxes F in the time series are divided by the mean flux. Histograms of $F_{band}/\langle F_{band} \rangle$ are given in Fig. 6. The distributions of the fluxes in each band are broadly consistent with being lognormal. Comparing the best-fit lognormal distributions to the data using the Kolmogorov-Smirnov test gives $P = 0.047$, $P = 0.091$, $P = 0.063$, and $P = 0.217$ in bands 1, 2, 3, and 4 respectively.

Power and lag spectra³, shown in Fig. 7, were generated for the broad

³The estimation of power and lag spectra was carried out principally by L. Miller

energy bands using the maximum likelihood method of Miller et al. (2010b) with $\Delta \ln \nu = 0.15$ (Legg et al. (2012), Sections 1.4 and 1.5).

3.1.2. 1H 0707-475




Fig. 8.— Histograms of $F_{band} / \langle F_{band} \rangle$ for each of the four broad energy bands of 1H 0707-495.

The 2008 observation of the narrow-line Seyfert I galaxy 1H 0707-495 (redshift $z = 0.0411$, Hewitt and Burbidge (1991)) has been the subject of numerous papers and vigorous debate. One major claim resulting from this object is the detection of emission from within $1R_g$ of the event horizon of the black hole (Fabian et al. 2009;




Fig. 9.— (Right) Power spectra of 1H 0707-495 in Band 1 (Black) and Band 4 (Grey), plotted as $\nu P(\nu)$. (Left) Lag spectrum between Band 1 and Band 4, positive values indicating that Band 4 lags behind Band 1.

Zoghbi et al. 2010), based on fitting of gravitational redshift to the claimed Fe-K and Fe-L lines. There has also been extensive examination of the lag spectrum (Zoghbi et al. 2010; Miller et al. 2010a; Zoghbi et al. 2011), with the discussion exemplifying the difference between the propagating fluctuations and hard band reverberation models. The mass has been estimated through use of the Kaspi et al. (2000) relation to be $M \approx 2.3 \times 10^6 M_{\odot}$ (Zhou and Wang 2005). However it has been noted that such a mass would mean that the object is radiating at ≈ 10 times the Eddington limit (Miller and Turner 2011; Leighly and Moore 2004).

The observations were made using the Large Window mode of the EPIC pn

detector, and the data used here is from the data reduction of Miller et al. (2010a). Source events were selected with $\text{PATTERN} \leq 4$ from a circular region $35''$ in radius. The background region was ‘about a factor of 5 larger’, and the threshold used for removing periods of high flaring background was $0.25 \text{ count s}^{-1}$. Miller et al. (2010a) report the data to be free of pile-up.

The mean fluxes in each of the four bands are 2.65 s^{-1} , 0.55 s^{-1} , 0.10 s^{-1} and 0.035 s^{-1} for Bands 1, 2, 3, and 4 respectively. Due to the low flux compared to Ark 564, the photons were binned with a temporal resolution of 256 s. For convenience and ease of comparison between bands the fluxes F in the time series are divided by the mean flux. Histograms of $F_{\text{band}} / \langle F_{\text{band}} \rangle$ are given in Fig. 8. The distributions of the fluxes in each band are consistent with being lognormal. Comparing the best-fit lognormal distributions to the data using the Kolmogorov-Smirnov test gives $P = 0.39$, $P = 0.73$, $P = 0.60$, and $P = 0.30$ in bands 1, 2, 3, and 4 respectively.

Power and lag spectra, shown in Fig. 9, were generated for the broad energy bands using the maximum likelihood method of Miller et al. (2010b) with $\Delta \ln \nu = 0.15$ (Miller et al. 2010a).

3.1.3. NGC 4051

An observation by Suzaku in 2008 found the Seyfert I AGN NGC 4051 (redshift $z = 0.0024$, Brinkmann et al. (1995)) in an unabsorbed state (Miller et al. 2010b). The Miller et al. (2010b) analysis showed that the variability of NGC 4051

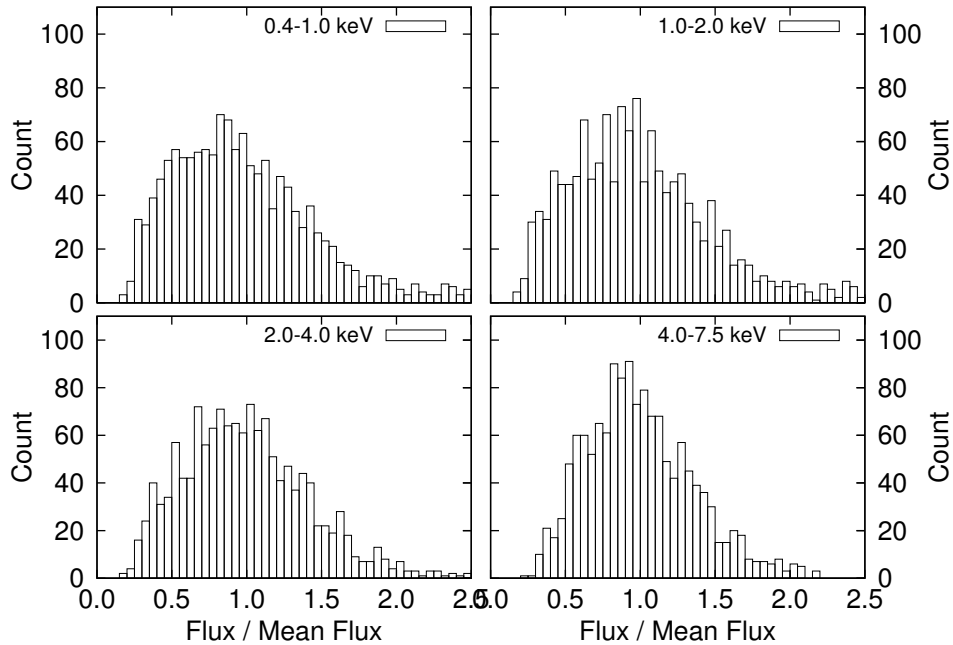


Fig. 10.— Histograms of $F_{band} / \langle F_{band} \rangle$ for each of the four broad energy bands of NGC 4051.

could be broken into two components - a soft unabsorbed component and a hard heavily absorbed component. The data was interpreted as due to partial covering of the source by clouds of absorbing gas. Early observations of the lag spectrum of NGC 4051 were quite crude, being split into just four frequency bins by Miller et al. (2010b). More recent analyses by Legg et al. (2012) and Alston et al. (2013) have improved the frequency resolution, with Alston et al. (2013) claiming to observe switching between positive and negative lags at low frequencies. Optical reverberation mapping by Denney et al. (2009) found the black hole mass to be $1.73_{-0.52}^{+0.55} \times 10^6 M_{\odot}$.

The 2008 dataset consists of 341 ks of observation time binned with a time




Fig. 11.— (Right) Power spectra of NGC 4051 in Band 1 (Black) and Band 4 (Grey), plotted as $\nu P(\nu)$. (Left) Lag spectrum between Band 1 and Band 4, positive values indicating that Band 4 lags behind Band 1.

resolution of 256 s. The mean fluxes in each of the four bands are 0.78 s^{-1} , 1.77 s^{-1} , 0.79 s^{-1} and 0.48 s^{-1} for Bands 1, 2, 3, and 4 respectively. For convenience and ease of comparison between bands the fluxes F in the time series are divided by the mean flux. Histograms of $F_{band}/\langle F_{band} \rangle$ are given in Fig. 10. The distributions of the fluxes in each band are consistent with being lognormal. Comparing the best-fit lognormal distributions to the data using the Kolmogorov-Smirnov test gives $P = 0.43$, $P = 0.07$, $P = 0.06$, and $P = 0.17$ in bands 1, 2, 3, and 4 respectively.

Power and lag spectra, shown in Fig. 11, were generated for the broad

energy bands using the maximum-likelihood method of Miller et al. (2010b) with $\Delta \ln \nu = 0.15$ (Legg et al. 2012).

3.2. Peak Selection and Measurement Parameters

The parameters used to obtain the peak responses are shown in Table 1. The shapes of the peak responses are relatively insensitive to changes in the parameters. The time series for AGN 1H 0707–495 and NGC 4051 were binned in intervals of 256 s owing to the noisiness of the data. Ark 564 has better quality data, and was binned at 128 s. Reducing the size of the time bins results in noisier peak responses, and difficulties in obtaining clear PSD and lag spectra (Section 3.1). Increasing the bin size reduces the resolution available on the peak response. The bin sizes were chosen to provide a good compromise between these two factors. The values chosen over-sample the peak by a factor 2-3, while the delayed excess is on a much longer timescale, and thus well sampled.

Since the results of the algorithm are only sensitive to timescales shorter than

Table 1: AGN peak detection and measurement parameters

AGN	Ark 564	1H 0707–495	NGC 4051
Sampling time (s)	128	256	256
Peak detection filter window length N	21	11	21
Peak detection filter length (s)	2688	2816	5376
Peak measurement filter window length N	81	41	81
Peak measurement filter length (s)	10368	10496	20736
Smoothing function FWHM (s)	370	530	500
Peak significance cutoff	2σ	2σ	2σ
Number of peaks found	117	65	79

the filter window length, the first-stage filter length was chosen to be sufficiently long that modes at frequencies below the break in the PSD were included. The third-stage filter length was a factor four larger, such that the low frequency modes that show positive lags in the lag spectrum were included. The similarities in the power spectra of Ark 564 and 1H 0707–495, and a desire for ease of comparison, meant that the filter windows for these objects were kept similar. In the case of NGC 4051 it was found that the filter lengths chosen resulted in significant negative residuals appearing in the final mean peak response, which disappeared when the filter lengths were increased by a factor 2, so in this case that longer filter length was adopted.

The main other effect of increasing the filter length is to increase the number of peaks detected, and to increase the significance of the hard band excess. The former effect is because as the filter length increases, the variations about the baseline level increase and more peaks meet the significance criteria. The latter effect likely arises because more low-frequency, positive-lag modes are included as the filter length is increased.

Due to concern about the possible impact of overlapping neighbouring peaks, a visual inspection of the located peaks was performed. While some peaks do overlap, the overall impact is symmetric - for every peak with another closely following it, there is a peak with one closely preceding. The final impact of the few overlapping peaks is minimal.




Fig. 12.— The extracted peak response for each of the three AGN: (top) Ark 564 sampled at 128 s; (centre) 1H 0707–495 sampled at 256 s; (bottom) NGC 4051 sampled at 256 s. Shaded regions indicate the 68% confidence region for each plotted point. Responses are shown for the 0.4–1 keV (dark shading) and 4–7.5 keV (light shading) bands. The peak responses have been normalized so that the amplitude at $t = 0$ is 1.0.

3.3. The Peak Shapes

The resulting peak responses are shown in Fig. 12. The most obvious feature is the coincidence of the soft and hard band peaks. Also of note is the asymmetry

of the hard band response in all three objects. While the asymmetry is most obvious in Ark 564 it is also present in 1H 0707–495 and NGC 4051. To examine the nature of the hard band asymmetry in Ark 564 the area under the central peak can be compared to the area under the delayed excess. The area of the peak region ($-448 \text{ s} < t < 448 \text{ s}$) is $4.55 \pm 0.48 \text{ s}$. The area of the excess is found by finding the difference between the region of the excess ($576 \text{ s} < t < 2496 \text{ s}$) and the equivalent region before the peak ($-2496 \text{ s} < t < -576 \text{ s}$). This evaluates to $8.55 \pm 2.31 \text{ s}$, indicating that there are 1.87 ± 0.55 times as many photons in the delayed excess as in the primary peak. In NGC 4051 and 1H 0707–495 the delayed excess and the primary peak cannot be easily separated, making it difficult to evaluate the fraction of delayed photons. Visual inspection indicates that those fractions are smaller in these two AGN compared with Ark 564.

3.4. Peak PSD analysis

It is important to verify that the peak responses extracted are a good estimate of the true response function. While testing with synthetic data has verified that the algorithm works, there is still the open question of how valid the underlying assumptions are. To do this, the peak responses shown in Fig. 12 are Fourier transformed, their power and lag spectra found and compared to those from the full time series. If the time series is a simple, stationary convolution of an underlying process with a response function, then it is expected that the spectra should have similar features.

To find power and lag spectra, the peak responses were padded to a length of

1001 time bins (i.e. much longer than the timescale of interest to avoid aliasing problems), and Fourier transformed. The Fourier modes were then binned to equal sized bins in log-frequency. The bin widths were sufficiently wide to eliminate correlation between adjacent modes. The low frequency information created in the padded series below the lowest measured values was discarded. The padding was important for eliminating aliasing caused by the periodic boundary condition of the Fourier transform, and prevents spurious modes being created within the frequency range of interest.

Bootstrap estimates were made of the statistical uncertainties on the Fourier modes (Efron 1979). For an averaged peak response formed from a weighted average of set of n individual peaks, the bootstrap method involved generating 1000 realizations, each consisting of n peaks drawn at random, but with replacement, from the observed set. In a given realization, some peaks could be present multiple times, others not at all. For each realization, peak responses and Fourier transforms were found and the variance of the realizations was taken to be the variance on the Fourier transform of the actual peak response.

The resulting PSDs are shown in Fig. 13, compared with the PSDs of the raw time series for the same observations. Because the peak response function has been renormalized, and because by selecting the highest peaks the amplitude cannot be considered typical of the entire time series, the peak shape PSDs have been normalized so that the total variance matches that of the full time series variance. The power spectra from the peak responses (Fig. 13) are in good agreement with those from the full time series in NGC 4051 and at low frequencies




Fig. 13.— Power spectra from the extracted peak responses of the three AGN: (left) Ark 564; (centre) 1H 0707–495; (right) NGC 4051. The top three panels display the soft (0.4–1 keV) band power spectra, while the bottom three panels show the hard (4–7.5 keV) band power spectra. Each panel compares the power spectrum of the peak response extracted (black circles) with the PSD of the full time series (grey squares). Vertical bars indicate the 68% confidence region on each point.

in the other AGN, but lack high-frequency power in Ark 564 and 1H 0707–495.

The break frequency in the soft band appears to be the same for the peak response as for the full series, although the power at frequencies above the break appears to be suppressed in the peak shapes. Three possible explanations of this discrepancy are:

1. There is an uncertainty in the position of the peaks detected in the smoothed time series. This uncertainty broadens the peak responses, suppressing high frequency power.

2. There might be additional high-frequency variations in the full time series, not captured by the peak selection.
3. The existence of an rms-flux relation (e.g. McHardy 2010) may cause aliasing of modes such that high-frequency power is boosted with respect to that expected in the simple ‘moving average’ model ⁴.

However, the good agreement between the PSDs at the frequencies where the positive lags are seen in all three AGN indicates that the peaks analysis procedure does capture valuable information, and that the peak responses found can be used as representative of the full series.

Lag spectra between the hard (4 – 7.5 keV) and soft (0.4 – 1 keV) bands were derived from the peak responses, and are compared with lag spectra from the full time series in Fig. 14. Uncertainties were estimated using the bootstrap process described above. The peak response lag spectra are more noisy, as expected considering that only a subset of the data has been used, but reproduce well the principal features at low frequencies, including the transition to lags around zero at high frequency. With the exception of the sharp negative lag feature in Ark 564,

⁴An rms-flux relation may be reproduced mathematically by taking a Gaussian process time series and multiplying by a function in the time domain. The Fourier transform of the resulting time series is then expected to be the convolution of the Fourier transform of the Gaussian process with the Fourier transform of the multiplying function. With a power-law power spectrum, such a convolution aliases power to higher frequencies.

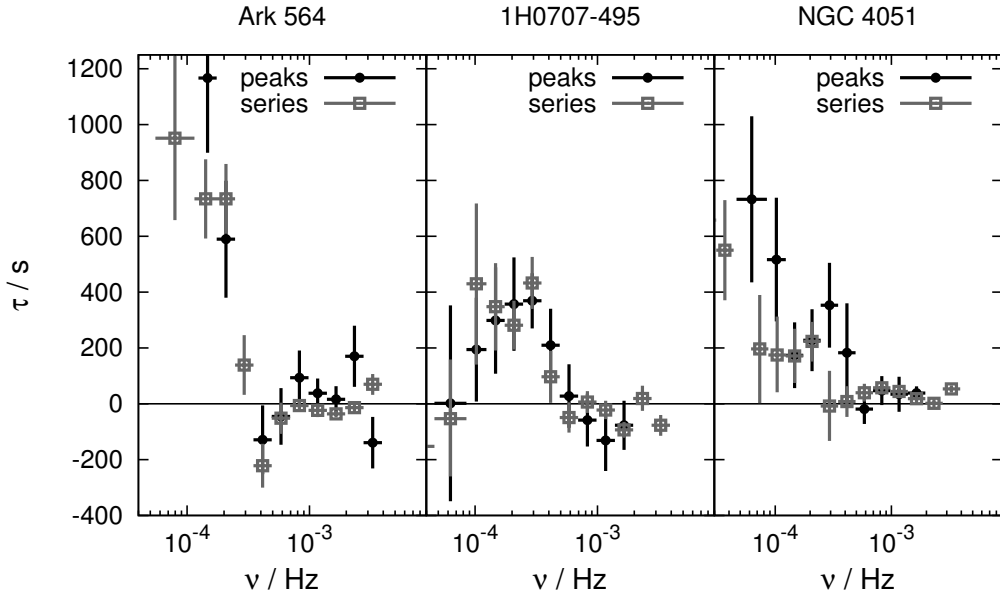


Fig. 14.— Lag spectra between 0.4 – 1 keV and 4 – 7.5 keV bands from the extracted peak responses (black points) compared with that derived from the full time series (grey points) for the three AGN: (left) Ark 564; (centre) 1H 0707–495; (right) NGC 4051. Vertical bars indicate the 68% confidence region on each point.

the transfer functions are too noisy to be able to reproduce the negative lags at high frequencies.

Overall, it can be concluded, from the similarity of both the PSDs and the lag spectra of the extracted peaks and the full time series, that the peaks analysis method does capture both the power spectrum and lag information that is obtained from a standard Fourier analysis, but now with the benefit of being represented as the time-domain response function in an individual energy band.




Fig. 15.— Lag spectra between each of the soft (0.4 – 1 keV, black) and hard (4 – 7.5 keV, grey) band peak responses and a δ function for the three AGN: (left) Ark 564; (centre) 1H 0707–495; (right) NGC 4051. Vertical bars indicate the 68% confidence region on each point.

As a test to distinguish between the competing models, lag spectra were calculated between the peak response in each band and a δ function (Fig. 15). This gives a measure of the asymmetry in each of peak responses, and is an important test for distinguishing between the proposed physical models. Under the ‘propagating fluctuations’ model it is expected that there exists an asymmetry in each band (Arévalo and Uttley 2006), and that each energy band would have a different, non-zero, lag with respect to a reference δ function. The observed lag

spectrum between two bands arises from the difference between the individual lag spectra. Under a reverberation model it is possible, albeit not guaranteed, that an individual band might have zero lag with respect to the reference δ function.

The results in Fig. 15 are hard to interpret for the object NGC 4051, but are more informative in the case of 1H 0707–495 and Ark 564. In 1H 0707–495 the soft band lag is at low frequencies, but not to a significant degree, and there are no signs of the short timescale soft band reverberation claimed by Zoghbi et al. (2011) as being responsible for the negative lag at high frequencies between the soft and hard bands. However, the proposed soft band reverberation lag is small, and is likely below the time resolution available with this analysis. In the case of Ark 564 the asymmetry is almost entirely in the hard band, with the soft band lag consistent with zero. This would seem to rule out the propagating fluctuations model used by Zoghbi et al. (2011) as an explanation for the X-ray time series observed in Ark 564. It does not, however, rule out alternative propagating fluctuation models, such as that of Gardner and Done (2014), which do not put such strict constraints on the role of each energy band in forming the lag spectrum.

3.5. The dependence of peak selection on energy band

Considering the extreme appearance of the hard band excess in Ark 564, further analysis was performed to examine its nature. To do this, the time series was investigated in a medium energy band (1–2 keV). When the peak selection is performed in this medium band, the resulting set of peaks is different from those selected in the soft band. This results in three sets of peaks - 52 detected only




Fig. 16.— Peak responses in soft and hard bands for three different selections of peaks in Ark 564: (top) peaks detected in both soft (0.4–1 keV) and medium (1–2 keV) bands; (centre) peaks detected in only the soft band; (lower) peaks detected in only the medium band. Shaded regions indicate the 68% confidence region for each plotted point. Responses are shown for the 0.4–1 keV (dark shading) and 4–7.5 keV (light shading) bands.

in the soft band, 29 detected only in the medium band, and 65 detected in both. These peak responses are shown in Fig. 16. The hard band peak response for this latter category shows a hard excess similar to that found previously. The set of peaks found only in the soft band produces an extreme hard excess, with a small primary peak. The set of peaks found only in the medium band shows no sign of a hard excess. A close examination of the peak positions shows that while no peak

was detected in the soft band, the positions are those of maxima in the soft band, just not significant enough to make the cut for selection. Whether this is due to noise or a systematic weakening of a subset of peaks in the soft band is unknown.

3.6. The effects of non-stationarity on the peak response in Ark 564




Fig. 17.— Peak responses in soft and hard bands for the first four observations (top) and the last four observations (bottom) of Ark 564. Shaded regions indicate the 68% confidence region for each plotted point. Responses are shown for the 0.4–1 keV (dark shading) and 4–7.5 keV (light shading) bands.

The underlying processes in Ark 564 have been shown to be non-stationary (Legg et al. 2012). In light of this, the peak responses obtained from each of the two halves of the dataset were examined. The resulting peak responses, shown in

Fig. 17, exhibit clear differences. Both peak responses exhibit an excess in the hard band at later times than the primary peak. The response from the first half of the data set has a larger excess, peaking with an intensity comparable with that of the primary peak and extending over ~ 2 ks. A crude comparison of the hard band primary peak with the excess gives an area for the primary of 3.9 ± 0.7 s over a range in time of ± 448 s, while the excess has an area of 10.9 ± 1.1 s over a range in delay time of 900–3000s. This corresponds to the hard band excess containing about 2.8 ± 0.6 times as many photons as the hard band primary peak. The response from the second half of the data set has a far smaller hard excess extending over just ~ 1 ks. The approximate comparison of the hard band primary peak with the excess gives an area for the primary of 5.5 ± 0.6 s while the excess has an area of 3.5 ± 0.7 s over a range in time of 1500–2700s, corresponding to the hard band excess containing approximately 0.6 ± 0.15 as many photons as the primary peak (the time ranges were chosen following visual evaluation of the period having the most significant delayed excess).

It is unclear how this effect and that shown in section 3.5 are linked. Table 2 lists the number of peaks detected in each of the observation periods. In the first four observations, peaks detected only in the soft (0.4 – 1 keV) band outnumber those detected in both soft and medium (1 – 2 keV) bands. In the last four observations the converse is true. Peaks detected only in the medium band appear to be uniformly spread. Thus the primary cause of the apparent variations in the delayed excess is not clear. Does the excess depend primarily on the peak spectral hardness ratio, with peaks of differing hardness ratios appearing in the time series preferentially at differing epochs? Or are both the variations in the delayed excess

Table 2: The number of peaks found in each observation, denoted by the OBSID, broken down by detection only in the soft (0.4 – 1 keV) band, in both soft and medium (1 – 2 keV) bands, or only in the medium band.

OBSID	0.4 – 1 keV	both 0.4 – 1 keV	1 – 2 keV
	alone	and 1 – 2 keV	alone
0670130201	7	9	3
0670130301	7	7	5
0670130401	8	4	5
0670130501	13	8	4
0670130601	5	9	3
0670130701	0	11	4
0670130801	6	8	2
0670130901	6	9	3

and the variations in the hardness of the detected peaks separate manifestations of the non-stationary variations in the source?

3.7. The nature of the delayed excess in Ark 564

The hard excess that is shown in Fig. 12 has been obtained by averaging many peaks together. The resulting mean peak responses for the hard excess could be caused by all peaks having a temporally broad, delayed signal with a hard spectrum. Alternatively, it might be that an individual primary flare tends to be followed by a narrow, secondary peak at a random time offset τ , where τ follows




Fig. 18.— The hard band (4 – 7.5 keV) response of Ark 564. The weighted mean response is shown in black with grey confidence region, while the weighted median is in red.

a probability distribution such that when averaged, we obtain the broad shapes seen. If the latter explanation were true, then it is expected that the excess in any given time bin would be dominated by a few individual peaks. To test for this, the weighted median of the individual peaks for each time bin is calculated, rather than a weighted mean. If the distribution of offset secondary peaks were the explanation, then the weighted median would fall below the weighted mean. In fact Fig. 18 shows that the weighted median agrees with the mean estimate within the uncertainties, indicating that the delayed excess in Ark 564 is indeed a continuous function which is roughly consistent between all flares.

3.8. Explanations for the extracted peak responses

3.8.1. *Intrinsic flare temporal structure*

At present, there is little understanding of the mechanism which creates the flares and the general time series variations. It is possible therefore that any temporal structure may be intrinsic to some complex high-energy process and not a consequence of propagation of either photons or sound through an extended distribution. For example, solar flares are generated by a rather poorly-understood mechanism, and also show features such as systematic spectral variations in hard X-rays (e.g. Krucker and Lin 2002). Obviously solar flares are complex and cannot be expected to be comparable with those produced by the very different environs of a supermassive black hole’s accretion disk. While the exact physical processes may differ, it does serve to illustrate the complexity and difficulty of the problem at hand.

Until the primary flare mechanism has been definitively established in AGN, we would never be able to eliminate this possibility. However, it can be noted that the timescale of the lags in Ark 564 is substantially longer than the duration of the primary flare, and it does seem likely that the delayed, hard-spectrum excess seen in Ark 564 arises through secondary reprocessing rather than through the primary flare process.

The durations of the primary flares measured in section 3.3 of 370 – 500 s allow upper limits to be deduced for the sizes of the regions producing those flares. The black hole masses are uncertain (section 3.1), but the flare durations imply the

flare generating regions have light-crossing times $28 r_g$ for Ark 564, $5 - 46 r_g$ for 1H 0707–495 (depending on the mass estimate used), and $57 r_g$ for NGC 4051.

3.8.2. Inverse Compton upscattering time delays

This analysis disfavours models which explain the lag as being due to inverse Compton upscattering of photons in the primary flare region, in which the hard band variations are delayed by the time taken to upscatter photons from the soft band. There is no sign of such a delay of the hard band peak with respect to the soft band peak in the peak shapes extracted: the positive lags appear to be entirely due to the additional hard-spectrum delayed component alone. The primary flares may be produced by inverse Compton upscattering, if the upscattering timescale is shorter than the temporal resolution of our analysis. The hard band excess has timescales a factor > 10 longer, not consistent with arising as inverse Compton upscattering delay in the flare generation region.

3.8.3. Positive lags from propagating fluctuations

Arévalo and Uttley (2006) have proposed that the low-frequency, positive lags in AGN arise from fluctuations in the accretion disk that propagate radially inwards, with an assumption that a fluctuation at large radius results in a soft X-ray emission spectrum, whereas a fluctuation at smaller radius results in a harder X-ray spectrum. This assumption is made because the accretion disk at small radii is hotter than at larger radii in the standard thin-disk model (Shakura and Syunyaev

1973). This is the explanation favoured by Zoghbi et al. (2011). The viability of such a model in AGN is open to question. It requires inward-propagating modes to dominate the variability of X-rays. Also, an AGN accretion disk itself is, unlike that of a galactic black hole candidate, not hot enough to emit significant hard X-rays.

Instead, the disk is required to act as a source of seed UV emission to an inverse Compton upscattering process, yet the process is thought to operate in the regime where the spectral index of the upscattered radiation is largely independent of the seed photon distribution (Titarchuk and Hua 1997). The power-law index of the Comptonized photons is instead dependent on the electron temperature and optical depth of the Comptonizing region. It is possible to get around this potential issue by allowing radiation from different parts of the disk to illuminate different Comptonizing regions with different properties, as in the alternative propagating fluctuations model proposed by Gardner and Done (2014). For instance it is not out of the question to imagine a black-body emitting disk with a coronal skin, with the corona electron temperature or optical depth varying with radius.

This model is also disfavoured by the measurements of Ark 564 presented here: the flaring behaviour shows sharp peaks that are temporally coincident at all photon energies, accompanied by a delayed, hard-spectrum signal, whereas in the simplest ‘propagating fluctuations’ model we would expect to see a broad temporal profile in which the spectrum hardened with time as the fluctuations propagate into the hotter regions of the accretion disk. Such a model might be consistent with the other two AGN, whose delayed signals appear as asymmetries in the peak

shape rather than in the form of a distinct component, but if we wish to find a single explanation for the lags in all AGN, the observation of time-coincident peaks accompanied by a distinct, hard, secondary maximum in Ark 564 remains challenging.

3.8.4. Energy Dependent Reverberation

The final model considered here is the reverberation explanation of Miller et al. (2010a,b) and Miller and Turner (2011). In those papers it was shown how a long-timescale delay could naturally produce low-frequency positive lags, and also potentially generate high-frequency negative lags as a result of the shape of the transfer functions in the two bands being cross-correlated (Miller et al. 2010a; Miller and Turner 2011). Such a delay might be produced by reverberation from gas a few tens to hundreds of gravitational radii from the X-ray source. Such models are consistent with the analysis presented here, in particular with the delayed hard excesses which are seen in all three AGN, but most prominently in Ark 564.

While the flare response in Ark 564 is characteristic of energy dependent reverberation, the strength of the hard band excess is a challenge to explain. The high fraction of scattered light requires a strong energy dependence, high covering fractions of the sky as seen by the source and/or high line-of-sight covering. This problem is examined further in Section 4.4.1.

4. Energy Spectral Analysis of Ark 564

Having applied the peak extraction algorithm of Section 2 to three AGN, the results in Section 3 indicate that the most interesting of the AGN is Ark 564. It is this AGN which exhibits the clearest difference in the peak response between soft and hard energy bands. The hard band peak response is strikingly similar in appearance to the toy top-hat reverberation models proposed by Miller et al. (2010a). This similarity, as well as the inadequacies of the alternative models (Section 3.8), leads to a preference for the reverberation explanation for the X-ray behaviour of Ark 564 and perhaps AGN in general.

This chapter further examines the object Ark 564 under the assumption of a reverberation model. This is achieved through exploring the data at a higher energy resolution. As well as the typical flux-energy spectra, lag-energy spectra can be formed from the full dataset of photon detection events. In Section 4.1 flux-energy spectral components associated with the peak and the delayed excess are extracted. In Section 4.2 full lag-energy spectra of Ark 564 are generated. Section 4.3 models the lag-energy spectrum associated with reverberation of the spectral components, and further demonstrates that a reverberation model is a valid interpretation for Ark 564. Section 4.4 discusses the physical interpretation of the flux-energy spectral components.

4.1. The Peak Energy Spectrum

The algorithm of Section 2.3 identified a set of times at which the soft (0.4 – 1 keV) band flux of Ark 564 peaked. These peak times were used to extract average flare responses in each of four energy bands, shown in Fig. 19. The analysis found 117 peaks in the soft band of Ark 564, and this analysis aims to examine them at a higher energy resolution than a mere four broad bands.

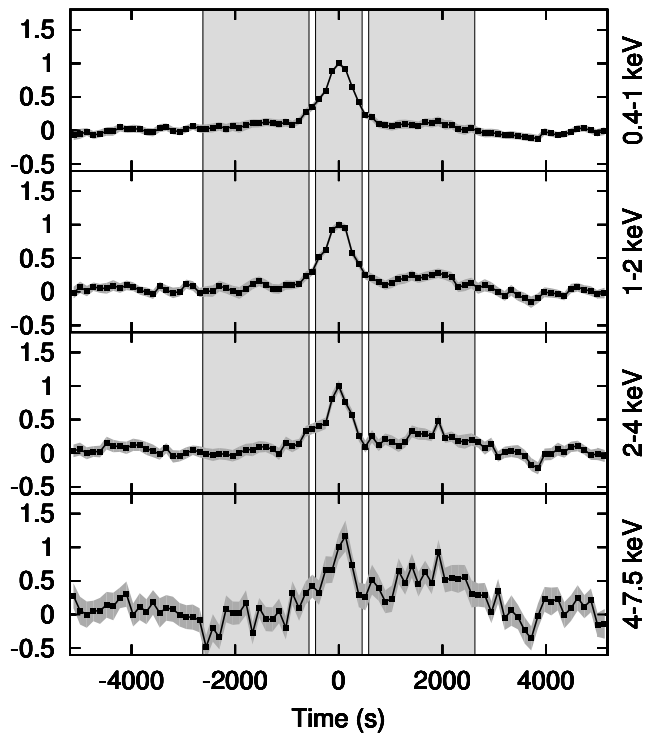


Fig. 19.— The flare responses extracted in Section 3.3 in each of the four broad energy bands (0.4-1 keV, 1-2 keV, 2-4 keV, 4-7.5 keV). 1σ confidence regions are in dark grey. The time bins used for the ‘before’, ‘peak’ and ‘after’ spectra are shaded in light grey.

The Ark 564 events list was binned into time bins of 128 s. For each of these time bins an energy spectrum was generated, with energy binning chosen to be the half-width at half maximum of the instrumental energy resolution, as in the formulation of Miller et al. (2007). These time resolved spectra were then used to generate three spectra, which are denoted as ‘before’, ‘peak’ and ‘after’. The ‘before’ spectrum was formed from 16 time bins ($-2624s < t < -576s$) before the identified time of each peak. The ‘peak’ spectrum was formed from 7 time bins ($-448s < t < 448s$) about each peak. The ‘after’ spectrum was formed from 16 time bins ($576s < t < 2624s$) after each peak.

There is inevitably some overlap between adjacent peaks, which means that the same bin may be used more than once. While this is cause for some concern, the overall effect is symmetric, and will not bias the resultant spectra too much. As was noted in Section 2.3, for every peak that that has another closely after it, there is a peak that has one closely before it. Time bins which had no valid data, such as occurred when a peak was detected near the end of an observation, were not included in the analysis.

The same peak weights as used by the peak response algorithm of Section 2.3 for the shape averages were used to create averaged spectra. The weights are given by the inverse of the shot-noise variance of the flux in the region around the peak ($-5184s < t < 5184s$) in the 0.4–1.0 keV band. The use of a unified weighting scheme allows for comparison with the results from the peak response analysis. The time intervals used to generate the three spectra are shown on Fig. 19.

The three spectra could not be immediately compared to the flare responses




Fig. 20.— The ‘before’ (light grey), ‘flare’ (dark grey, crosses) and ‘delayed’ (black, dots) spectra. For clarity, the spectra have been rebinned using the XSPEC `ipl` command `setplot rebin 10 2`. Error bars are 1σ confidence regions obtained through bootstrapping. For reasons of clarity and space the ‘peak’ or ‘after’ spectra are not shown, but are replaced by their residuals after subtracting the ‘before’ spectrum.

due to the baseline subtraction that was performed to obtain the flare responses. The assumption was made that the three spectra shared a common base spectrum, which would correspond to the baseline variations removed by the peak response algorithm. Of the three spectra, the ‘before’ spectrum was taken to come closest to this baseline, being largely uncontaminated by the peak or delayed excess. The maximum amount of the ‘before’ spectrum was subtracted from the ‘after’

spectrum, while maintaining positivity of the residual. The maximum amount to subtract was arrived at by dividing the ‘after’ spectrum by the ‘before’ spectrum, and finding the energy bin with the minimum ratio. Subtracting more than this amount would result in energy bins with a negative value, which is clearly unphysical. Since the presence of higher degrees of noise at high energies could throw off such a division, the search for the minimum ratio was restricted to the 20 energy bins $E < 1.85$ keV. As both the ‘peak’ and ‘after’ spectra are harder than the ‘before’ spectrum, this restriction does not cause any bins to take a negative value. The same amount of the ‘before’ spectrum was also subtracted from the ‘peak’ spectrum, resulting in the spectra seen in Fig. 20. The residual of the ‘peak’ spectrum is denoted as the ‘flare’ spectrum, and the residual of the ‘after’ spectrum as the ‘delayed’ spectrum.

This maximal subtraction of the ‘before’ spectrum is only one possibility, chosen to illustrate the differences in shape between the three spectra. The ‘true’ spectra associated with the flares and the delayed component could lie between these spectra and the full, unsubtracted spectra. It is to be noted, though, that the peak responses require a large difference between the ‘delayed’ and ‘flare’ spectra, with the ‘delayed’ spectrum expected to be harder. The difference between the two spectra can be maximised by removing as much as possible of the common component.

To form uncertainty estimates for these spectra, bootstrap resampling was used (Efron 1979). The bootstrap sampling method avoids the difficulties of propagating shot noise errors through a lengthy analysis, and implicitly takes into

account any correlations in the data. It achieves this by sampling the population, and examining the variance of the samples. Each sample was made up from 117 peaks drawn at random from the original list of 117 peaks, with the possibility of duplication. The full analysis was performed for each sample. The 1σ errors were taken to be the standard deviation of an ensemble of 200 such samples. The resulting error margins are larger than those shown by Giustini et al. (2015) due to the bootstrap sampling capturing variations in spectral shape.

It can be seen from Fig. 20 that the ‘delayed’ component is indeed harder than the ‘flare’ component, as should be expected from the peak responses which exhibit a delayed excess that increases with energy band (Fig. 19). It is also apparent that the spectral components are variations around a common baseline spectrum (the ‘before’ spectrum) of order 10%. These extracted spectra are further examined and discussed in Sections 4.3 and 4.4.

4.2. Lag-Energy Spectra

A common tool for the examination of X-ray time series in AGN is the lag-energy spectrum. The lag-energy spectrum examines the phase lag between the flux in a narrow energy band and the flux at all other energies. The most common method of plotting this is to plot the lag in a specific frequency range as a function of energy (Zoghbi et al. 2011; Kara et al. 2013). The approach of Zoghbi et al. (2011) was followed using the more rigorous Fourier analysis methods of Miller et al. (2010a).

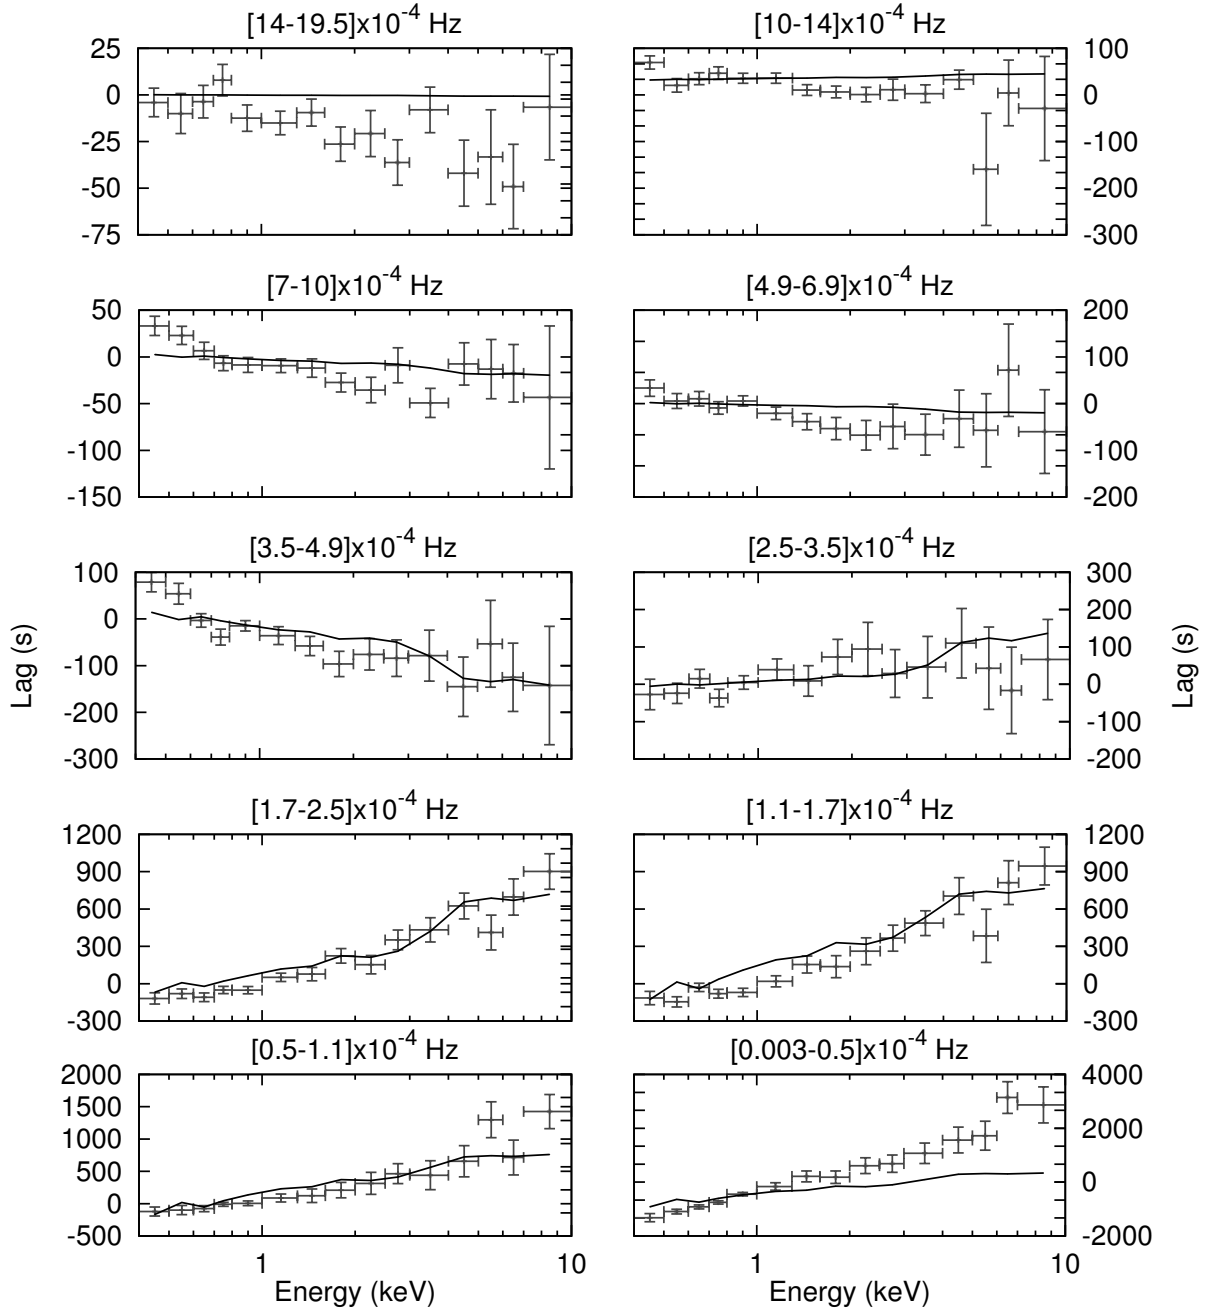


Fig. 21.— Lag-energy spectra in ten frequency bands. The observed lag-energy spectra are in grey, while the model spectra are given by the black lines. Note the varying scales of the y-axis. Error bars mark the 68% confidence regions of the observed lags and the widths of the energy bins.

For each energy bin in the lag-energy spectrum, time series $S_E(t)$ and $H_E(t)$ are constructed. The energy bins were chosen to be the same as in Kara et al. (2013). $H_E(t)$ consists of photons in that energy bin, while $S_E(t)$ consists of all other photons in the range 0.4 – 10 keV. The two time series do not share any photons, which avoids correlated noise. The lag-energy spectrum for a given frequency consists of the average lag in that frequency range plotted against energy.

To calculate the required lag spectra the maximum-likelihood method of Miller et al. (2010a) was used. This method simultaneously fits frequency-binned power and lag spectra to the data, and is able to account for gaps in the data. Once the maximum likelihood parameters are found through gradient ascent, confidence regions are estimated by stepping in each parameter until the $\Delta\chi^2 = 1$, while other parameters are allowed to vary. This method produces robust error estimates and allows for correlated uncertainties, and differs from the method used by Kara et al. (2013) to generate lag-energy spectra. The method has since been tested and verified by Zoghbi et al. (2013).

Lag-energy spectra for all frequency bins are shown in Fig. 21. It is to be noted that this analysis does not reproduce the peaked lag-energy spectrum found in Fig. 3 of Kara et al. (2013) in the $[3.2 - 5.2] \times 10^{-4}$ Hz frequency range. The apparent feature in Kara et al. (2013) was interpreted as corresponding to the Fe K line at $\tilde{6.4}$ keV, and was used to support a conclusion of short timescale reverberation as part of the ‘propagating fluctuations’ model. However, the lags are correlated between energy bins, and there is no computationally practical method of evaluating the magnitude of the correlations. This means that the actual

significance of an apparent feature in three adjacent energy bins is much lower than it might at first appear. This analysis uses a different time binning (128s compared to 10s for Kara et al. (2013)) and allows for gaps in the time series due to the removal of flares in background emission (Legg et al. (2012), section 2.1). The fact that the apparent feature is not reproduced by this analysis indicates that it was most likely a statistical anomaly.

4.3. Modelling the Lag-Energy Spectrum

In Section 4.1 a breakdown of the energy spectrum was extracted, and interpreted in a simple model as being formed of two components - a ‘flare’ spectrum, and a reverberating delayed excess. If this is a valid interpretation, then it should agree with the lag-energy spectra of Section 4.2. To test this, the most basic of the reverberation models was applied - a response function consisting of a δ function at $t = 0$ with a delayed top-hat of width Δt centred at time t_0 . The delayed top-hat integrates to a fraction f of the delta function. This is the simplest model proposed by Miller et al. (2010a), and while it cannot capture the full complexity of lag spectra, it does provide a decent first-order fit and is easy to conceptualize. The top-hat model is not intended as a realistic physical model (arising from a thin spherical shell with a hole towards the observer), but as a mathematical construct that explains the observations.

In applying this simple model to the extracted spectra, the delta function is associated with the ‘flare’ spectrum and the top-hat with the ‘delayed’ spectrum. The top-hat fraction f is thus energy dependent, and is set as the ratio of the flux

in the ‘flare’ and ‘delayed’ spectra. The values of Δt and t_0 are assumed constant across all energy bins.

This is used to build up model lag-energy spectra for given values of Δt and t_0 , and by comparing the model to the true lag-energy spectra the maximum-likelihood values were found to be $\Delta t = 1702$ s and $t_0 = 1633$ s. The model lag-energy spectra are shown as solid lines in Fig. 21. The fit is good apart from at the very highest and lowest frequencies. The difference at high frequencies could be due to statistical fluctuations, especially when we consider the likelihood of correlations between energy bins. The low frequency difference is more dramatic, and is due to the overly basic nature of the top-hat response function.

Further insight into the nature of the lag-energy spectrum can be gained by noting the equation derived by Kotov et al. (2001).

$$\tau_\nu(E) = a_\nu \log E + b_\nu \tag{6}$$

This equation arises naturally from a model in which the lag spectrum is caused by variability in the power-law index, and arises due to the fact that the ratio between two power-laws is linear in $\log E$. In the reverberation picture, that ratio is directly associated with the ratio between the δ function and the area of the delayed response function. Fitting such an equation to the data results in $\chi^2 = 132.8$ for 130 degrees-of-freedom. Plotting the gradient a_ν against ν (Fig. 22) gives a shape reminiscent of that seen between broad soft and hard energy bands (Section 3.1.1). The lag gradient has been overlaid with three examples of the lag spectrum from the fitted top-hat model, with values of $f = (1.0, 0.5, 0.2)$. From this we can see that while a simple top-hat model is a decent fit at mid to high frequencies, it does




Fig. 22.— (Left axis, black errorbars) Gradients of the log-linear fits to the lag-energy spectra. Error bars indicate frequency bin widths and 95% confidence regions. (Right axis, grey lines) Lag spectrum associated with the fitted top-hat response function for three values of f : $f = 1.0$ (upper, dashed), $f = 0.5$ (middle, dotted), $f = 0.2$ (lower, dash-dot).

not provide a large enough lag at low frequencies.

Both the semi-empirical log-linear model and the top-hat reverberation model exhibit divergence from the measured lag spectra in the lowest frequency bin, on time scales longer than 20 ks. This divergence cannot be readily explained with simple models, and is outside the scope of the peak response analysis of Section 3.

These long time scales span the gaps between the eight individual observations, and will be affected by the non-stationarity observed in the system.

4.4. Physical Interpretation

The extraction of a change in spectral shape associated with the peak responses of Section 3 confirms the results of the previous analysis with different methods at a higher energy resolution. The interesting result comes from modelling the lag-energy spectra associated with the simple reverberation model, and comparing it to the observed lag-energy spectra.

The quality of the fit indicates that this reverberation picture is a valid possible explanation of the data worthy of further consideration. The most obvious criticism is that the proposed top-hat reverberation model does not explain the lowest frequencies. This criticism, that top-hat models do not fully explain every detail of lag-spectra, has been a persistent and recurring one since they were proposed by Miller et al. (2010a). What such criticisms miss is that the top-hat models are purposefully simplistic, intended to demonstrate that a reverberation model can produce lag spectra of the required form. Indeed, Section 3 demonstrates that the extracted flare responses, reminiscent of a modified top-hat, have associated lag spectra consistent with that of the full time series.

Now that spectra associated with the direct and delayed components of a reverberation model have been obtained, the task is to model those spectra to gain an insight into the physical nature of the reverberation. Due to the noisy nature of

the spectra there is no expectation of obtaining a detailed and unique spectral fit, and in fact all the fits performed are over-fitting the data with a minimal number of model components. What can be expected is to be able to put limits on possible models, and use those limits to constrain the physical properties of the system.

4.4.1. Absorption Model




Fig. 23.— The ‘flare’ (black) and ‘delayed’ (red) spectra unfolded against the XSPEC model `tbabs*(po*mtable + atable)`. The `mtable` and `atable` are absorption and emission tables generated by XSTAR. Error bars indicate 1σ confidence regions obtained by bootstrapping.

Qualitatively, the ‘delayed’ spectrum in Fig. 20 appears as if it could be an absorbed version of the ‘flare’ spectrum. The most that can be expected is

to be able to put limits on the nature of the absorber. To this end, XSPEC was used to model the ‘flare’ spectrum as a simple power-law with galactic absorption (`tbabs*power`), while the ‘delayed’ spectrum is modelled with the same power-law and galactic absorber, with the addition of an `mtable` absorber and `atable` emission. The table components of the model come from a grid of XSTAR tables spanning column densities between $N_H = 1 \times 10^{18} \text{ cm}^{-2}$ and $N_H = 3.16 \times 10^{24} \text{ cm}^{-2}$, and ionizations in the range $0 < \log \xi < 5$.

The best fit slab absorber found was a low ionisation ($\log \xi \approx 0.0$) absorber with column $N_H = 5.1 \pm 1.6 \times 10^{21} \text{ cm}^{-2}$, fully covering the delayed component (Fig. 23). The XSTAR definition of $\xi = L_{0.0136-13.6 \text{ keV}} / \rho r^2$ (erg cm s^{-1}) was used, where ρ is the number density of ionised material, and r its distance from the source with luminosity L . The fit is a clear case of overfitting, with XSPEC reporting a reduced chi-squared of 0.64. The result can, however, be used to put rough limits on the physical nature of the reverberation, using a model of a cold Compton scattering reflector with a cold absorber. From the ratio of the delayed and direct components it is required that the reflected component have a magnitude roughly equal to that of the direct. This presents difficulties when it comes to constructing a physical model.

The first issue is that it requires a very high covering fraction of the source by the scatterer. This can be alleviated to an extent by the presence of high column ‘bricks’ that suppress the direct component by completely blocking a fraction f_b of the direct component without imprinting any absorption features. A diagram showing the principal physical components is shown in Fig.24. The




Fig. 24.— Diagram illustrating the required physical components for an absorption model. The source of the photons emits with luminosity L_0 in all directions. The direct component is blocked by high column ‘brick’-like absorbers with covering fraction f_b . The delayed component arises from a scatterer with covering fraction f_s which imprints absorption on the delayed spectrum.

spectral components extracted imply that at high energies the direct and scattered components have roughly the same luminosity. The direct component has an intensity $(1 - f_b)L_0 \frac{\delta\Omega}{2\pi}$, where f_b is the covering fraction of the high column ‘bricks’, L_0 is the luminosity of the source which is assumed to radiate uniformly in the hemisphere above the accretion disk, and $\delta\Omega$ is the solid angle covered by the observing detector. The delayed component has intensity $f_s L_0 \frac{\delta\Omega}{4\pi}$, where f_s is the covering fraction of the scattering material. The scattering material is assumed to scatter uniformly over a full 4π solid angle.

$$(1 - f_b)L_0 \frac{\delta\Omega}{2\pi} \simeq f_s L_0 \frac{\delta\Omega}{4\pi} \quad (7)$$

$$f_b \simeq 1 - \frac{f_s}{2} \quad (8)$$

A useful rule of thumb is that a wind with an opening angle of 45° has a covering fraction $f_s = 0.7$. If this value is taken for f_s , then the ‘bricks’ must cover $f_b \simeq 0.65$ of the source. This is assuming that source luminosity is isotropic. If it instead originates from a slab with a $\cos\theta$ angular dependence the required covering fractions will be even higher.

The second issue is that of the absorber. The simplest way to add this into the picture would be as a thin slab covering the entirety of a cold Compton scatterer. The physical plausibility of this simple addition is, however, rather dubious. An alternative would be to embed the absorber within the scattering material. This would require high density contrasts. The ionisation of the bulk of the Compton scatterer cannot be below $\log \xi \approx 4$ without imprinting the absorption spectrum on the observations. The ionisation of the absorber cannot be above $\log \xi \approx 0.5$. Since the absorber and scatterer are at roughly the same distance from the source, the absorber must have a density $> 10^{3.5}$ times that of the scatterer. When the low column depth ($\approx 5 \times 10^{-3} \times 10^{24}$) of the best fit absorber is accounted for, it is apparent that the length scale of the absorber is $\approx 10^{-6}$ of the length scale of the Compton scattering. This again presents issues for physical plausibility.

4.4.2. Comptonization Model

As an alternative to finding a physically plausible absorption model, one in which the ‘delayed’ spectrum arises from inverse Compton scattering by hot electrons is also considered. To test this, both the ‘flare’ and ‘delayed’ spectra were modelled with a galactic absorber and the Comptonization model of Titarchuk (1994) (`tbabs*comptt`). The `comptt` model assumes a black-body source of seed photons with temperature T_0 at the centre of either a sphere or slab of gas with optical depth τ_p and temperature T_p . This simple model is intended to simply fit the apparent power-law nature of each component.




Fig. 25.— The ‘flare’ (black) and ‘delayed’ (red) spectra unfolded against two XSPEC Comptonization models. (Left) `tbabs*comptt`, (Right) `tbabs*simpl*comptt`. Error bars indicate 1σ confidence regions obtained by bootstrapping.

This is easily able to model the data (Fig. 25, left), with a reported reduced chi-squared of 0.5. The parameters of the model are given in Table 3. Exploratory fits indicate that T_p and τ_p are inversely correlated. The low value of χ^2 is a clear indicator that even with these simple models the data are being overfitted, though the true significance is lower due to the correlations between energy bins.

Table 3: Comptonization model parameters

Spectrum	Galactic Column 10^{22} cm^{-2}	T_0 keV	T_p keV	τ_p	norm
‘Flare’	0.0543	$8.4 \pm 4.6 \times 10^{-2}$	55 ± 191	0.10 ± 0.22	$0.33 \pm 1.3 \times 10^{-4}$
‘Delayed’	0.0543	0.21 ± 0.039	20 ± 329	1.1 ± 19.7	$0.47 \pm 7.8 \times 10^{-5}$

The parameters for the model `tbabs*comptt`. The input black body temperature is given by T_0 , the temperature of the plasma electrons by T_p and the optical depth of the plasma by τ_p .

The Comptonization model has plenty of scope for variation. While the ‘flare’ spectrum is fitted with `comptt`, it could equally well be fitted with a simple power-law or the more complex `optxagn` (Done et al. 2012) model. The degeneracy between plasma temperature and optical depth gives scope for different physical geometries.

Unfortunately the `comptt` model restricts the user to using a black-body as input to the Comptonization. An alternative possibility is that the ‘delayed’ spectrum is an upscattered version of the ‘flare’ spectrum. This can be tested with the XSPEC model `simpl`, which is a convolution model that turns an

arbitrary input spectrum into a power-law at high energies, while maintaining the correct low-energy behaviour (Steiner et al. 2009). Unfortunately this model (`tbabs*simpl*comptt`) does not give any physical properties of the Comptonizing region, as it takes the index of the power-law as a parameter. The fact that it does produce a good fit ($\chi^2 = 83.21$, $dof = 168$, Fig. 25 (right)) does indicate that the ‘flare’ and ‘delayed’ components can be explained by the same source spectrum, with the ‘delayed’ component having been reprocessed by a hot Comptonizing region. The physical properties of the Comptonizing region can be estimated from the `comptt` fits, but are poorly defined.

The question which then needs to be answered is whether such a Comptonization model is physically plausible, and whether any assumptions made are appropriate. The `comptt` model can be switched between a disk and spherical geometry, neither of which are appropriate for the desired reverberation model, since in each case the source of the seed photons is interior to the Comptonizing gas (Sunyaev and Titarchuk 1985). The geometry envisaged is one in which the seed photons have an unimpeded path to the observer, while also illuminating the Comptonizing region, and one of the simplest ways to achieve this is an axisymmetric wind (Fig. 26) such as that arising from a displaced model for line driven winds (Sim et al. 2008; Knigge et al. 1995).

The angular dependence of the Compton emission from a plane is a linear function of $\mu = \cos \phi$, where ϕ is the angle with respect to the normal to the plane, and $F(\mu)$ can be well approximated as $F(\mu) = \frac{1}{2} + \mu$ (Sunyaev and Titarchuk 1985; Magdziarz and Zdziarski 1995). Considering an observer at angle θ to the axis of




Fig. 26.— Diagram illustrating the geometry of a simple wind. The source of the photons emits in all directions, and is observed at an angle θ to the axis of symmetry. The region containing the wind is between the disk and an inner face at an angle ψ to the disk. The directly observed flux is L_0 , while the delayed flux is L_d .

symmetry, and an inner wind edge at angle ψ to the disk, it can be found that the average value of μ over the full inner edge is $\bar{\mu} = \cos \psi \cos \theta$. The wind covers a fraction $f_c = \sin \psi$ of the sky as seen by the source. If isotropic illumination by the source is assumed, and that half of all photons entering the wind are backscattered out of it, then $L_d = f_c L_0 \frac{F(\bar{\mu})}{2}$. Applying the appropriate substitutions gives the following equation:

$$L_d = L_0 \times \frac{\sin \psi}{2} \left(\frac{1}{2} + \cos \psi \cos \theta \right) \quad (9)$$

If it is instead assumed that the source radiates in the manner of a disk, the situation becomes more complicated. The source itself has an angular dependence, $L_0 = \bar{L} \left(\frac{1}{2} + \cos \theta \right)$, where \bar{L} is the average source flux. For the delayed component it is necessary to average over the covered angles, and obtain the following equation:

$$L_d = \bar{L} \times \frac{1}{4} \left(\sin \psi + \frac{1 - \cos 2\psi}{2} \right) \left(\frac{1}{2} + \cos \psi \cos \theta \right) \quad (10)$$

When the photon counts given by the `tbabs*simpl*comptt` model are examined, it is found that the model source spectrum contains 5.8 times as many photons as the Comptonized component. This ratio could be achieved with a value of ψ as low as $\psi \simeq 15^\circ$ (for $\theta = 0.0$) in the isotropic case, and $\psi \simeq 30^\circ$ in the case of emission from a plane.

4.4.3. *The Preferred Scenario*

It has been demonstrated that a reverberation model is a valid interpretation of the X-ray data obtained for Ark 564. A top-hat response function is a surprisingly good fit to the data, and spectral components associated with the direct and delayed components of such a reverberation model have been extracted. Two possibilities for the physical cause of the change in spectral shape have been explored.

The first of these, in which the spectral changes are caused by absorption by the distant scatterer, is possible to achieve with spectral fitting. Such fits become implausible, however, when the required physical geometry is considered. The nature of absorption, in which photons are ‘lost’, means that extremely high covering fractions of the source by the absorber are required. Even then, additional obscuration of the direct light is required.

The second model brings about the change in spectral shape not by absorbing

photons but by Compton scattering them to higher energies. This shift of photons from low to high energies provides the required suppression at low energies, and boosts the higher energies. The result is that the required ratio of direct to scattered light is achieved at much lower covering fractions than required by the absorption model. The Comptonization model is preferred for the comparative ease with which the required scattered fractions can be achieved, in contrast to the extreme scenarios required for absorption models to work.

5. Angular Dependence of Comptonization of X-ray Photons

The modelling of simple Comptonization scenarios to explain the behaviour of Ark 564 (Section 4.4.2) revealed the shortfalls in existing modelling of Comptonization. Sunyaev and Titarchuk (1980) demonstrated that a generic feature of Comptonization by hot electrons is a power-law distribution to high energies, and Sunyaev and Titarchuk (1985) later showed the angular distribution of escaping Comptonized photons. The principal parameters for Comptonization are the plasma temperature kT_e and the optical depth τ_0 of the Comptonizing material. Hua and Titarchuk (1995) further explored the behaviour of Comptonization for a range of values of these two parameters, and restricted the regimes of applicability (Fig. 7, Hua and Titarchuk (1995)).

These studies were limited to exploring Comptonization where the source of the seed photons is interior to the Comptonizing material, whether it be an infinite slab or a sphere. There have also been studies of the Comptonization of an exterior source of photons. Gorecki and Wilczewski (1984), following on from the work of Pozdniakov et al. (1977), employed Monte-Carlo simulations to study the transmission of X-rays through slabs of gas. Magdziarz and Zdziarski (1995) explored the angular dependence of reflection from a slab of cold electrons both with and without absorption. The reflection was averaged over a range of incidence

angles, mixing together grazing incidence with photons travelling perpendicular to the slab.

In the light of the preferred model of Section 4.4.2, this work seeks to expand the modelling of Comptonization in two directions. Firstly, to examine the reflection of X-ray photons by hot Comptonizing material. Secondly, to explore the dependence of the resulting spectra and fluxes on the incidence angle of the seed photons. Taking advantage of advances in computing technology allows a deeper exploration of the parameter space than was available to previous studies.

5.1. Compton Scattering

The behaviour of photons upon scattering by photons is described by the Klein-Nishina formulae (Longair 1992; Klein and Nishina 1929), and is dependent on the energy of the incident photon. Upon scattering, the photon experiences a change in energy which is related to the change in direction of the photon

$$E' = \frac{E}{1 + \alpha(1 - \cos \theta)} \quad (11)$$

where E and E' are the photon energy before and after scattering respectively, $\alpha = \frac{E}{m_e c^2}$ is the ratio of the photon energy to the electron rest mass, and θ is the angle between the incident and scattered photon directions.

The angular distribution of the scattered photons is not uniform, except in the low photon energy limit. Instead, the angular distribution is given by the

Klein-Nishina Differential Cross-Section

$$\frac{d\sigma_{KN}}{d\Omega} = \frac{r_e^2}{2} \left(\frac{E'}{E} \right)^2 \left(\left(\frac{E'}{E} \right) + \left(\frac{E'}{E} \right)^{-1} - 1 + \cos^2 \theta \right) \quad (12)$$

where $\frac{d\sigma_{KN}}{d\Omega}$ is the cross section to scattering per infinitesimal solid angle, and r_e is the classical electron radius. For practical purposes, it is easiest to normalise by the Klein-Nishina integrated cross-section and use it to describe the probability distribution of scattering angles.

$$\sigma_{KT} = 2\pi r_e^2 \left(\frac{1 + \alpha}{\alpha^2} \left[\frac{2(1 + \alpha)}{1 + 2\alpha} - \frac{1}{\alpha} \ln(1 + 2\alpha) \right] + \frac{1}{2\alpha} \ln(1 + 2\alpha) - \frac{1 + 3\alpha}{(1 + 2\alpha)^2} \right) \quad (13)$$

These three equations can be used to describe the probability of a photon being scattered by a stationary electron, the probability distribution of the scattering angle and the change in energy. In practice, of course, electrons need not be stationary. A plasma will contain electrons with a distribution of velocities and directions of travel, and this motion must be accounted for. This is achieved by shifting to the rest frame of the electron, which travels with velocity $\beta = v/c$ in direction $\hat{r} = (r, \theta, \phi)$.

The first step is to rotate the coordinate system such that the electron travels in a directions parallel to the principal axis. This is achieved with the rotation

matrix A which operates on the photon's 4-vector $\hat{P} = (E, p_x, p_y, p_z)$.

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \sin \theta \cos \phi & \sin \theta \sin \phi & \cos \theta \\ 0 & \sin \phi & -\cos \phi & 0 \\ 0 & \cos \theta \cos \phi & \cos \theta \sin \phi & -\sin \theta \end{pmatrix}, \quad A^{-1} = A^T \quad (14)$$

The relativistic Doppler shift is then given by applying the Lorentz boost L .

$$L = \begin{pmatrix} \gamma & \beta\gamma & 0 & 0 \\ \beta\gamma & \gamma & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad L^{-1} = \begin{pmatrix} \gamma & -\beta\gamma & 0 & 0 \\ -\beta\gamma & \gamma & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (15)$$

The 4-vector P'_i of the incident photon in the electron rest frame is related to the 4-vector P_i of the photon in the simulation frame by applying these two transformations in succession.

$$\hat{P}'_i = LA\hat{P}_i \quad (16)$$

After performing the Compton scattering in the electron rest frame, the scattered photon has the inverse transforms applied to bring it back to the simulation reference frame.

$$\hat{P}_s = A^{-1}L^{-1}\hat{P}'_s \quad (17)$$

5.2. The Monte-Carlo Simulation

The Monte-Carlo simulation draws photons from a given incident spectrum and propagates them through a region of plasma, scattering from electrons where appropriate. Escaping photons are added to the output spectra.

One way of drawing photons is to use an input spectrum with a tidy analytic form, such as a power-law distribution. Such distributions can be rearranged and integrated to map a uniform distribution between 0 and 1 to the spectral probability distribution. This allows the use of random variables drawn from the pseudo-random number generators (PRNGs) provided by many numerical computing packages. This approach has its downsides though. While the resolution of drawn values is limited only by the resolution of the PRNG, the method relies on the desired distribution having a form that can be treated analytically, a very small subset of possible distributions. Common distributions such as the black-body distribution are not possible to rearrange and integrate into the desired form. It is also desirable to be able to make use of spectra generated by spectral fitting programs such as XSPEC.

The alternative approach is to use a table of values and weights. A photon can be drawn uniformly from the table of energies, and an appropriate weight assigned. As well as allowing for arbitrary distributions, this approach also improves the statistics of the simulation at high energies. Since AGN often have steep spectra, the low counts at high energies produce high shot-noise errors. The weighting approach avoids reproducing these shot-noise errors in the simulation since the actual counts of simulated photons are uniform with energy. A potential downside

is that energies are restricted to those listed in the table, but this can be alleviated through interpolation.

On the face of it, the propagation of photons through the plasma is a hard problem. To determine the distance a simulated photon travels before a scattering event requires a draw from a distribution governed by the mean-free-path λ of the photon. Naïvely, λ is governed by the average electron cross-section, which would require averaging the Klein-Nishina integrated cross-section over a range of energies due to Doppler shifting of the photon energy by the thermal distribution of electrons. This would be a computationally formidable task. Fortunately, the situation can be resolved by noting that the electron cross-section is in fact effectively invariant under Lorentz transformation (Melrose 2008). The argument is that the rate of Compton scattering R is governed by $R = \sigma cn_{ph}$, where σ is the electron cross-section and n_{ph} is the number density of photons. Since both R and n_{ph} transform as time, $\sigma = \frac{R}{cn_{ph}}$ is invariant under Lorentz transformation.

This invariance simplifies the problem. The Klein-Nishina integrated cross-section σ_{kn} can be calculated for the simulated photon, and used to draw a travel length from the distribution $P(x)dx \propto \exp(-\rho\sigma_{kn}x)dx$ given by the Beer-Lambert law. If the photon is still within the geometrical bounds of the gas distribution after travelling that distance, an electron velocity and direction are drawn from the appropriate thermal distribution. The photon then undergoes a Lorentz transformation to the electron rest frame, where the Compton scattering is performed, before the inverse transformation is applied to bring it back into the simulation frame.

To avoid the difficulties of drawing a direction for the scattered photon from the complicated angular distribution given by the Klein-Nishina differential cross-section, a new direction is drawn from a uniform spherical distribution. The photon is then weighted appropriately according to the differential cross section, and the energy shifted in accordance with Eq. 11.

The process can be summarised as follows:

1. Draw a photon from the desired energy distribution.
2. Calculate the Klein-Nishina integrated cross section for the photon energy (Eq. 13).
3. Draw a travel distance from the distribution $P(x)dx \propto \exp(-\rho\sigma_{kn}x)dx$.
4. If travelling that distance means the photon escapes the system, go to 13. Otherwise, go to 5.
5. Draw an electron velocity and direction from the appropriate distribution.
6. Apply the transformation $\hat{P}'_i = LA\hat{P}_i$ into the electron rest frame (Eq. 16).
7. Draw a direction of travel for the scattered photon from a spherical distribution.
8. Calculate the Klein-Nishina differential cross section (Eq. 12).
9. Weight the photon by the differential cross section.
10. Calculate and apply the energy shift (Eq. 11).

11. Apply the inverse transformation $\hat{P}_s = A^{-1}L^{-1}\hat{P}'_s$ back into the simulation frame (Eq. 17).
12. Go to 2.
13. Write the photon weight to the output array.

5.3. Verifying the Simulation Output

It was important to verify that the simulation does in fact produce a valid output. This was done by comparing to the output of the XSPEC models `comptt` and `nthcomp`. The `comptt` model is based on the work of Titarchuk (1994), and makes use of the Wien black-body approximation $n(E) \propto E^2 e^{E/k_B T}$ for the input spectrum. The model geometry consists of an infinite planar slab of thickness τ with photon sources situated along the centreline of the slab (Fig. 1a of Sunyaev and Titarchuk (1985)). The model takes the temperature of the black-body, the electron temperature of the slab, and the slab thickness as parameters. The `nthcomp` model is more empirical in nature, treating the Comptonized spectrum as a power-law of index Γ with low and high energy cutoffs given by the black-body and electron temperatures respectively.

The test parameters for the black-body and electron temperatures were chosen to be $k_B T = 0.1$ keV and $k_B T = 20$ keV respectively, and the slab thickness was chosen to a column of depth $\tau = 4\tau_0$, where τ_0 is the optical depth to Thompson scattering. The power-law index for the `nthcomp` model was to chosen to be $\Gamma = 1.9$. Two Monte-Carlo simulations were performed, one using a black-body

spectrum as the source spectrum, and the other using the Wien approximation. The Monte-Carlo simulations were performed with 4.95×10^8 photons each, and the results are shown in Fig. 27.

It can be seen from Fig. 27 that the high energy behaviour of the Monte-Carlo




Fig. 27.— Comparison between two Monte-Carlo simulations and two XSPEC models. The Monte-Carlo simulations consist of a slab with thickness $\tau = 4\tau_0$ and an electron plasma temperature of $k_B T = 20$ keV. The Monte-Carlo simulation input spectra are a black-body (light grey) and the Wien approximation (dark grey) respectively, both with temperature $k_B T = 0.1$ keV. The `comptt` model (blue) has an optical depth $\tau = 4\tau_0$, plasma temperature $k_B T = 20$ keV, and input temperature $k_B T = 0.1$ keV. The `nthcomp` model (red) has an power-law index $\Gamma = 1.9$, high temperature cutoff $k_B T = 20$ keV, and low temperature cutoff $k_B T = 0.1$ keV. The ratio between the two XSPEC models is also shown in black.

simulations is independent of the input spectrum, as expected. The low energy behaviour varies depending on whether the input spectrum is a true black-body or the Wien approximation. The high and low energy cutoffs agree with those predicted by the `nthcomp` model. The discrepancy lies in the comparison with the `comptt` model. At low energies, the model does not agree with the Wien approximation, while the high energy discrepancy appears to be due to a difference in the power-law index. This could be caused by either a difference in the electron distribution (which would be expected to cause a shift in the location of the high energy cutoff) or a difference in the electron cross-sections.

There is also a discrepancy between the `comptt` and `nthcomp` models. The difference in low energy cutoff is to be expected, since the two models use different input spectra. The difference in the power-law region is due simply to the `nthcomp` power-law index Γ being chosen to fit to the Monte-Carlo simulations. What is more of a problem is the discrepancy in the high energy cutoff. If this was simply an offset due to a difference in the power-law indices it would be expected that the ratio of the two models be constant in this region. Instead the two models have a different shape. The disagreement between the `comptt` and `nthcomp` models, as well as rigorous testing of the simulation components, leads to a tentative conclusion that the discrepancy is not due to an issue with the simulation. This discrepancy has been noted by Zdziarski et al. (1996), the originators of the `nthcomp` model.

The agreement of the Monte-Carlo simulations with the high and low energy cutoffs predicted by the `nthcomp` model leads to a conclusion that the disagreement between Monte-Carlo models and `comptt` is not as problematic as it might seem.

The `comptt` model was chosen because it is the only published spectral model which attempts to explain the physical properties of the Comptonizing region. Differences in assumptions made may give subtly different spectral results, but the greater flexibility of the Monte-Carlo approach still allows insights to be reached. That the results of simulation differ from the `comptt` model does indicate that they may not be suitable for detailed spectral fitting, but this is neither practical nor the aim of the simulation. The `nthcomp` model is just as, if not more, widely used (being a component of the `optxagn` model), and is in good agreement with the simulation.

5.4. Reflection from a hot, infinite slab of finite optical depth

The geometry of the reflection problem is similar to that used by the `comptt` model. A slab of finite thickness has infinite extent in the x-y plane, and is populated by electrons with temperature $k_B T_e$. The incident photons are initialised at the lower surface of the slab, and travel in the direction $\hat{r} = (1, \theta_i, \phi = \pi)$, where θ_i is the angle between the photon direction vector and z-axis. The geometry is shown in Fig. 28. The photons are drawn from a distribution generated with the `nthcomp` XSPEC model. The lower and upper model temperatures are 0.01 keV and 300 keV respectively, and the power-law index is $\Gamma = 2.0$.

The emergent photons are binned into 800 angular bins of equal solid angle. These cover 40 equal intervals between $\phi = 0$ and $\phi = 2\pi$, and 20 equal intervals between $\cos \theta = -1$ and $\cos \theta = 1$. Photons which pass through the slab have values $\cos \theta > 0$, while reflected photons have $\cos \theta < 0$. The normalised photon




Fig. 28.— Diagram representing the slab geometry. The slab is infinite in extent in the x-y plane, and has thickness τ along the z-axis. Incident photons P_i make an angle θ_i with the slab normal and travel parallel to the x-axis in the x-y plane. Outgoing photons P_o make an angle θ_o with the slab normal and ϕ_o with the x-axis.

counts are shown in Figures 29-32 for two slab thicknesses and a range of incidence angles.

The most readily obvious effect is that there can be a strong angular dependence on ϕ as well as $\cos\theta$. The dependence on $\cos\theta$ is well known from previous studies (e.g. Sunyaev and Titarchuk 1985), while the ϕ dependence becomes apparent due to the novel approach here of performing the simulation with a single incidence angle. On consideration, the result is not overly surprising. It is well known that the Klein-Nishina differential cross-section can have a strong

dependence on the change in direction of the photon, with photons more likely to be scattered forwards or backwards rather than perpendicular to the initial direction of travel. The strength of the dependence on ϕ varies with incidence angle, being strongest when the incident photon travels near parallel to the slab, due to this behaviour of the differential cross-section.

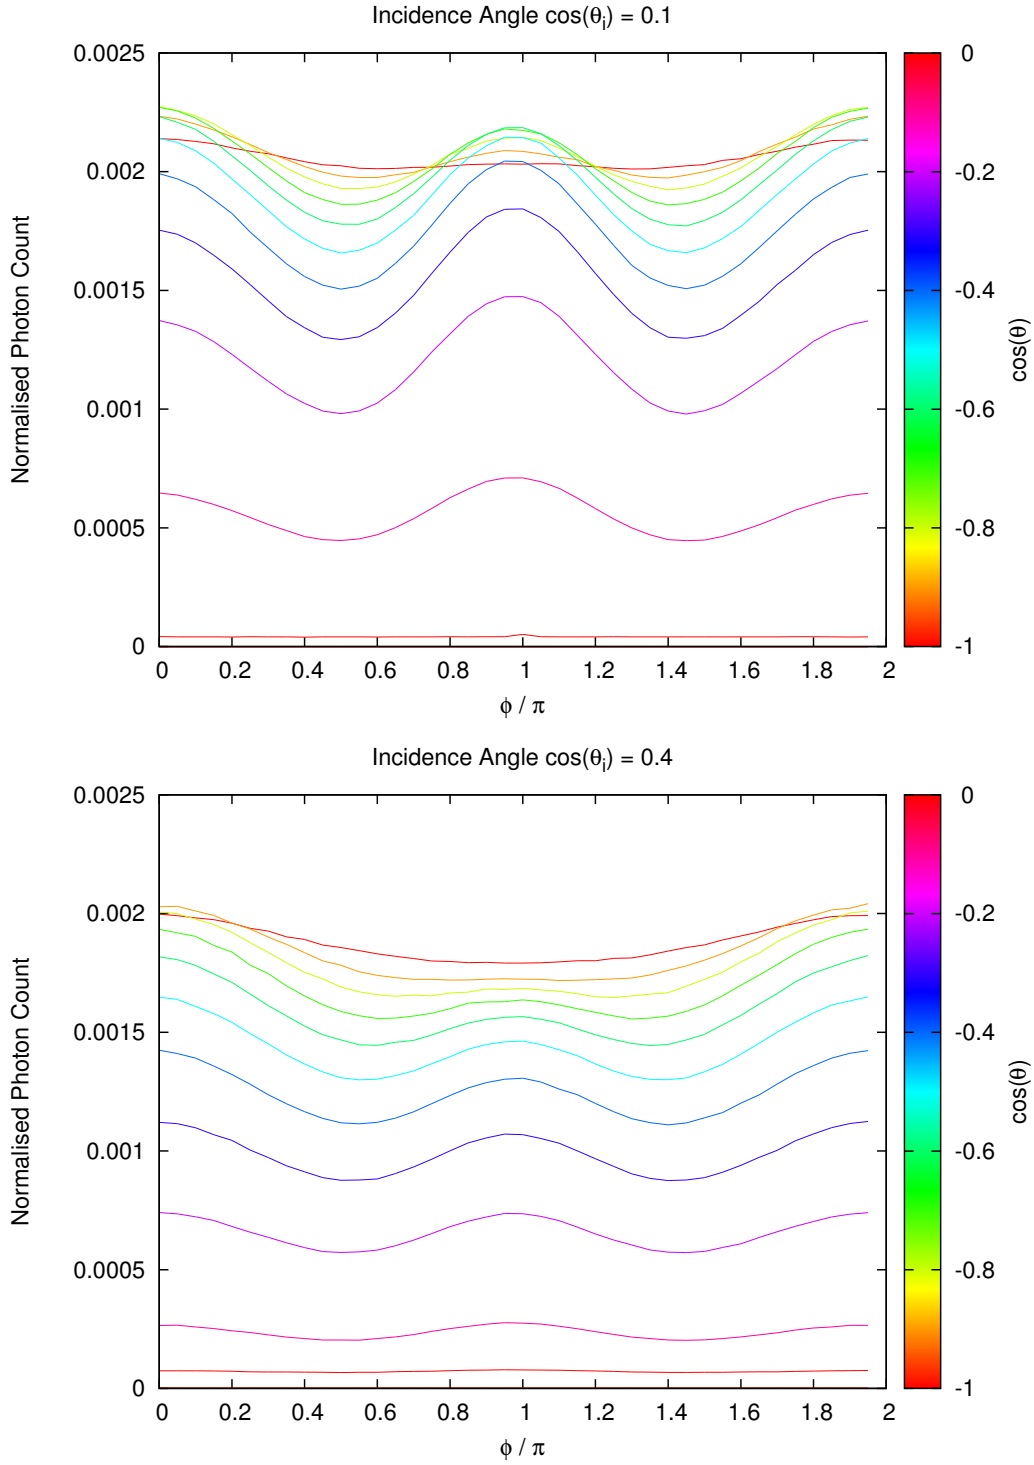


Fig. 29.— Reflection of photons with incidence angles $\cos \theta_i = 0.1$ (top) and $\cos \theta_i = 0.4$ (bottom) from an infinite slab of thickness τ_0 and temperature $k_B T_e = 20$ keV.

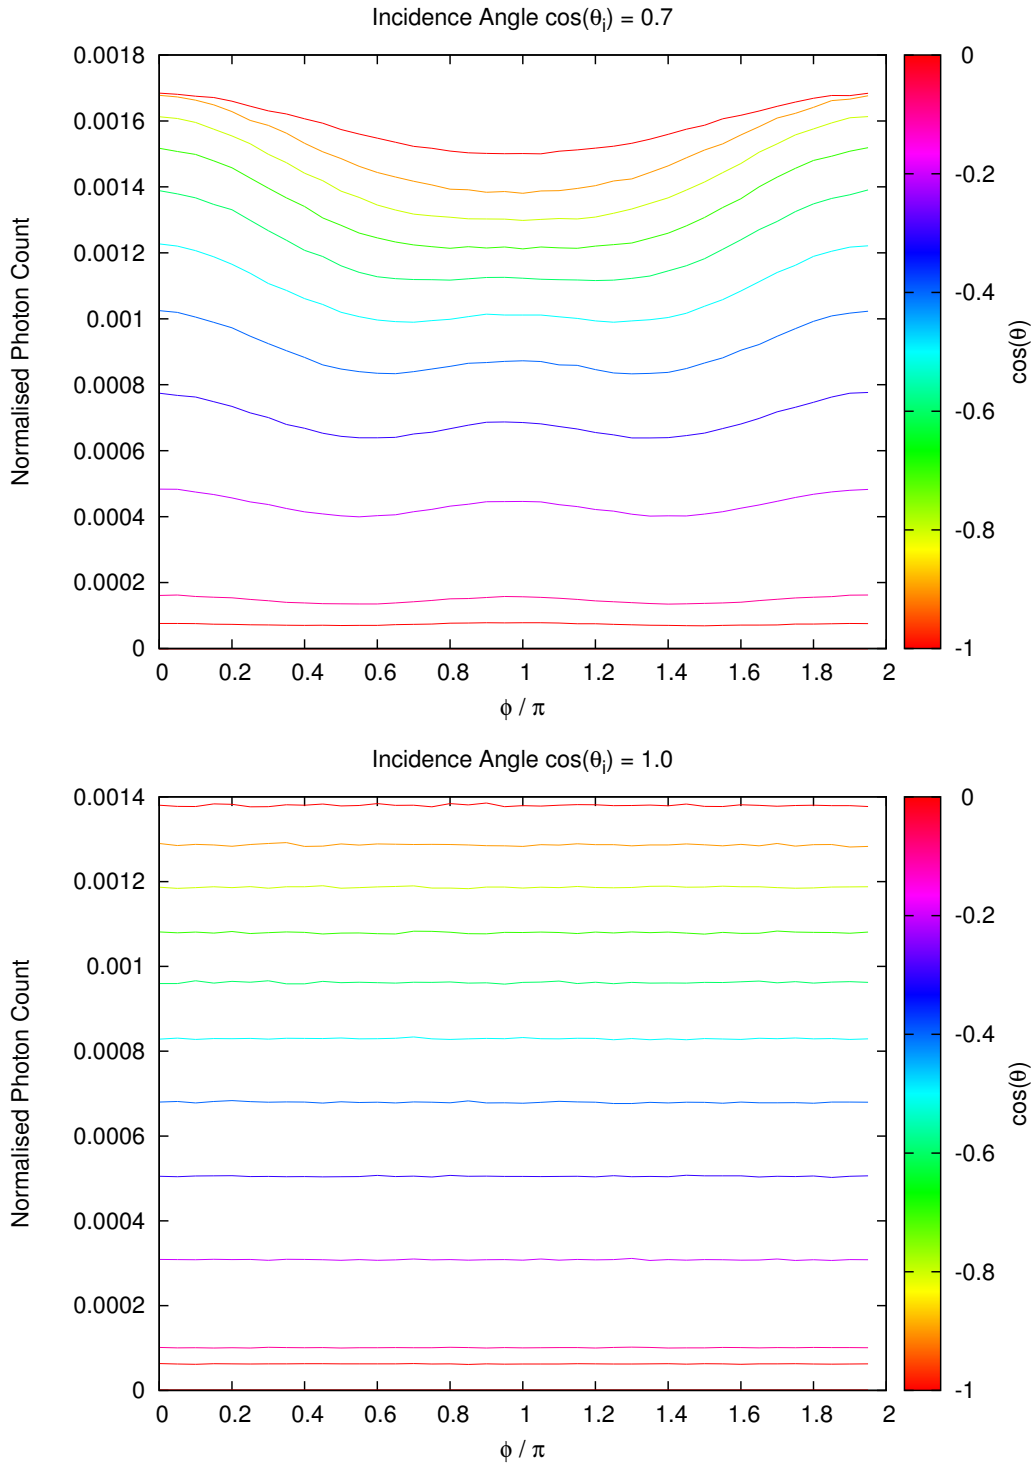


Fig. 30.— Reflection of photons with incidence angles $\cos \theta_i = 0.7$ (top) and $\cos \theta_i = 1.0$ (bottom) from an infinite slab of thickness τ_0 and temperature $k_B T_e = 20$ keV.

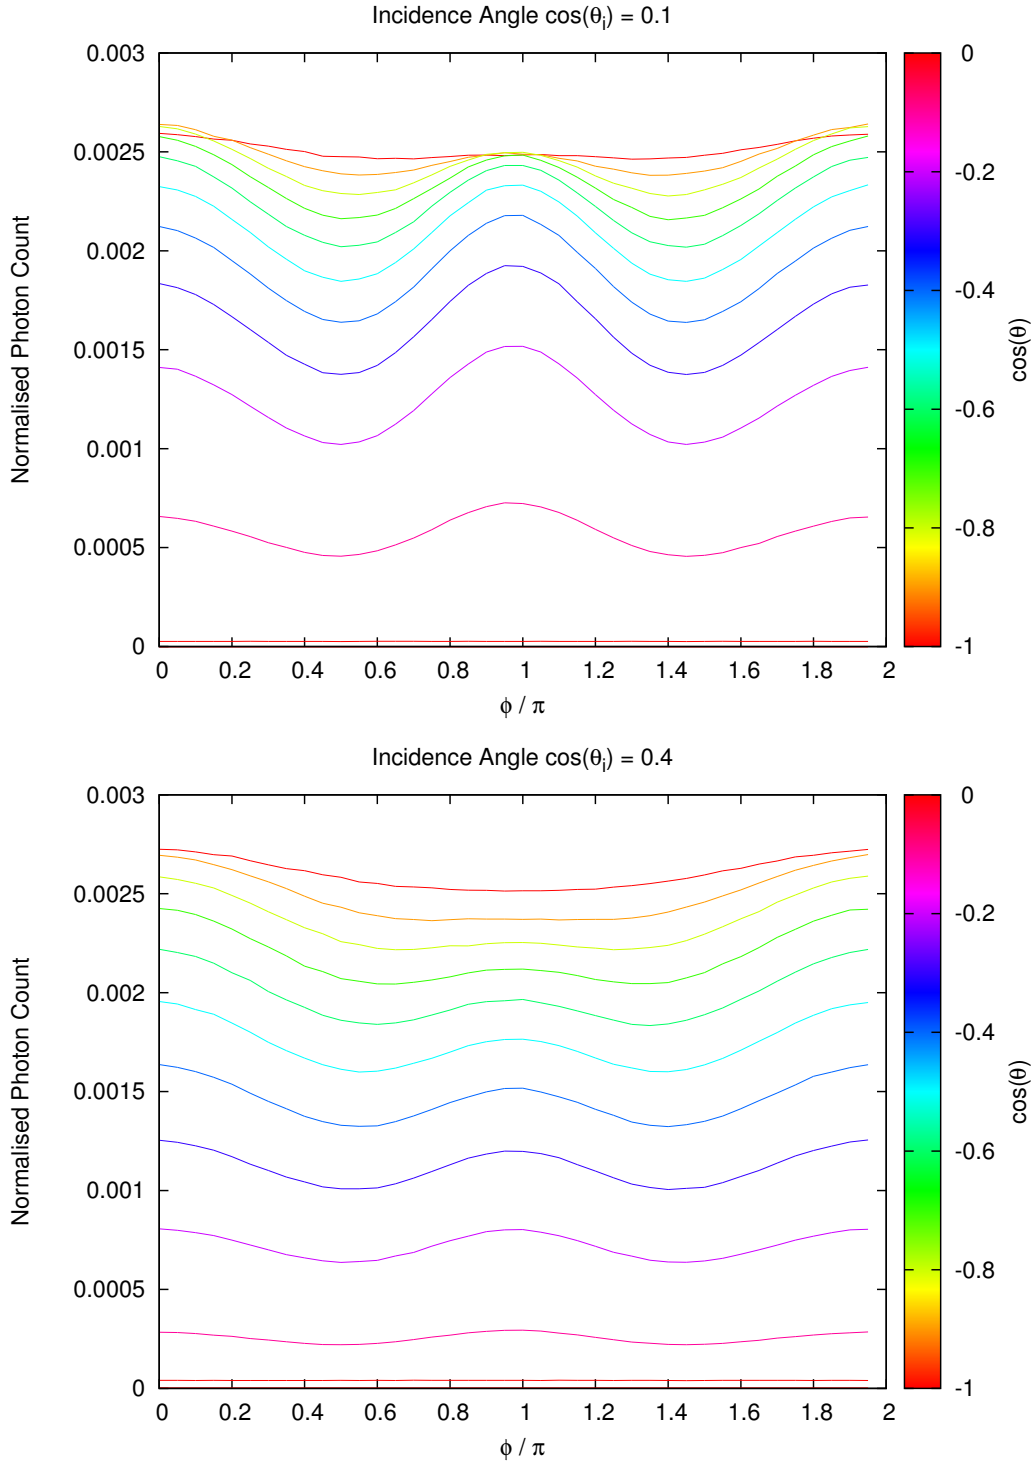


Fig. 31.— Reflection of photons with incidence angles $\cos \theta_i = 0.1$ (top) and $\cos \theta_i = 0.4$ (bottom) from an infinite slab of thickness $2\tau_0$ and temperature $k_B T_e = 20$ keV.




Fig. 32.— Reflection of photons with incidence angles $\cos \theta_i = 0.7$ (top) and $\cos \theta_i = 1.0$ (bottom) from an infinite slab of thickness $2\tau_0$ and temperature $k_B T_e = 20$ keV.

There is also a dependence on the optical depth of the slab. With an optical depth of $\tau = \tau_0$, half of the reflected photons have undergone just a single scattering. It is photons which undergo a single scattering that carry the imprint of the differential cross section, each successive scattering weakening the effect. The distribution of the number of scatterings undergone is shown in Fig. 33 (Top) for four different slab thicknesses. Increasing the slab thickness reduces the fraction of photons which undergo but a single scattering while increasing the portion which have undergone higher numbers of scatterings.

This variation in the number of scatterings undergone is what gives rise to the increasing degree of spectral hardening with rising slab depth (Fig.33, Bottom). Visual comparison with XSPEC models such as `comptt` indicates that the location of the high energy turn off depends solely on the slab temperature, and not on the model geometry or shape of the input spectrum. It is clear, however, that the other features such as the spectral index are highly dependent on the input spectrum as well as the slab depth and temperature.




Fig. 33.— (Top) Distribution of the number of scatterings undergone by photons reflected by slabs of optical depth $\tau = \tau_0, 2\tau_0, 4\tau_0, 8\tau_0$. (Bottom) Spectral output from each of the slabs with $k_B T = 20$ keV and an incidence angle of $\cos \theta = 0.7$. The input spectrum is shown in black. Output spectra for each of the optical depths are shown for $\cos \theta = 1.0$, each summed over all angles in ϕ .

5.5. Comparison to Ark 564

The purpose of this foray into Comptonization was to evaluate the possibility that the spectral components identified in Section 4.1 are due to Comptonization of the reverberating component. Unfortunately, it is not feasible to use this Monte-Carlo simulation to fit a model to the spectra of Fig. 20. This is due to the long run-time of the simulation, that the best-fit parameters are likely to be poorly constrained, and the noisy nature of the data. What is possible is to use simulation parameters which provide an output spectra that is qualitatively similar to the observed spectra.

The proposed model to explain the ‘flare’ and ‘delayed’ spectra identified in Section 4.1 is that the ‘flare’ spectrum is seen directly, while the ‘delayed’ spectrum is the result of Comptonizing reflection of the ‘flare’ spectrum. The input spectrum for the simulation is set as the `nthcomp` XSPEC model which well describes the ‘flare’ spectrum. This model has a black-body temperature of 0.1 keV, a power-law index of $\Gamma = 3$, and a high energy cutoff of 300 keV. The photons drawn from this distribution have an incidence angle of $\cos \theta_i = 0.7$. The slab has an optical depth of $\tau = 4\tau_0$ and a temperature of 20 keV. A total of 11.9×10^9 photons were generated using 120 computing cores with a runtime of six and a half hours.

In examining the output spectra, the variation in total flux with $\cos \theta$ must be accounted for. This is achieved by normalising the spectra such that they have a value of unity at 1 keV. As can be seen in Fig. 34, there is significant angular variation in the spectral form. The greatest degree of spectral hardening is seen at angles near perpendicular to the disk, while angles near parallel to the disk have




Fig. 34.— Spectral output from the simulation of a slab with optical depth $\tau = 4\tau_0$, $k_B T = 20$ keV, and an incidence angle of $\cos \theta = 0.7$. The input spectrum is shown in black. Output spectra are shown for each bin in $\cos \theta$, each summed over all angles in ϕ . The bins in $\cos \theta$ have upper limits in integer multiples of 0.1. Spectra are normalised to unity at 1 keV.

the least hardening. At low energies the flux is significantly reduced due to photons being shifted to higher energies.

It is possible to delve deeper into the simulation output, and examine the spectral variation with azimuthal angle ϕ . The simulation output is much noisier when resolving individual angular bins. Due to this, the spectra shown in Fig. 35 have been smoothed with a boxcar function 5 energy bins in width for improved




Fig. 35.— Spectral output from the simulation of a slab with optical depth $\tau = 4\tau_0$ and $k_B T = 20$ keV and an incidence angle of $\cos\theta_i = 0.7$. The input spectrum is shown in black. Output spectra are shown for each bin in $\cos\theta$. Spectra are shown for every other bin in ϕ with the forward direction shown in black and backscattering in light grey. Spectra are normalised to unity at 1 keV, and only every 11th point is shown.

clarity, and only every 11th point is shown. It can be seen that the spectra diverge, with the forward direction having a softer spectrum than the backward.

Like the other parameters related to the system geometry, the incidence angle




Fig. 36.— Spectral output from the simulation of a slab with optical depth $\tau = 4\tau_0$, $k_B T = 20$ keV, and an incidence angle of $\cos \theta_i = 0.1$. The input spectrum is shown in black. Output spectra are shown for each bin in $\cos \theta$, each summed over all angles in ϕ . The bins in $\cos \theta$ have upper limits in integer multiples of 0.1. Spectra are normalised to unity at 1 keV.

has a large impact. The results of a separate simulation with incidence angle set to $\cos \theta_i = 0.1$ are shown in Fig. 36. Comparison with Fig. 34 shows that the shallower incidence results in much softer output spectra.

This angular analysis shows that care should be taken when considering reflection scenarios and geometries, but also that there is a rich variety in the behaviour which can be seen from even simple models. Not only does the flux vary

significantly with viewing angle, but so too does the spectral shape. A single choice of slab parameters and incidence angle can look very different when viewed from different angles.

The implication of this for Ark 564 is that even simplistic models, such as the simple wind model suggested in Section 4.4.2 (Fig. 26), quickly become very complex to work with. Every point on the inner surface of such a wind would have a different reflected flux and spectral shape. Travelling around the inner surface azimuthally sketches out a locus in the $\theta - \phi$ plane of observation angles, the exact shape of which depends strongly on both the wind opening angle and disk inclination. As well as this, the incidence angle varies with distance from the central source. To properly evaluate such a wind model would require a new simulation in which the infinite slab is replaced by a wind geometry. Such a task is beyond the scope of this work.

5.6. Spherical reflectors

An alternative approach is to look at spherical reflectors rather than infinite slabs. Photons are incident from a single direction on a uniform sphere of gas with radius R . The cross-section of the sphere is πR^2 , and the optical depth τ varies as $\tau = (R - r)\tau_c$ where r is the distance from the centreline and τ_c is the optical depth along the centreline. The symmetries in this model eliminate the incidence angle θ_i and output angle ϕ as variables of interest. The model is also attractive due to its finite volume being more physically plausible than an infinite slab.

Explorations with this basic model show that the resulting spectra require high maximum optical depths and high temperatures to produce significantly hardened spectra. This is due to the nature of the geometry. To produce spectral hardening, photons must be boosted from low to higher energies. This requires a large number of scatterings, with the average energy boost per scattering dependent on the temperature. The spherical geometry means that the average optical depth of an incident photon is much lower than τ_c , and the finite size means that photons can more easily escape than in an infinite slab scenario. These factors conspire to make the number of scatterings undergone before escape lower than with an infinite slab, and so require higher temperatures to achieve a significant boost in energy.

An example of the output spectra is shown in Fig. 37. The sphere of gas has a maximum optical depth along the centreline of $\tau_c = 8\tau_0$ and a temperature of $kT = 100$ keV. The input spectrum was chosen to be the same as that of Section 5.5. The output in each of 100 evenly spaced bins in $\cos\theta$ are shown, and exhibit little difference in hardness. There is a trend toward harder spectra in the forward directions than the back-scattered, but the degree of variation is small.

It is possible to use this information on the angular distribution of reflection from a single sphere to consider possible geometries to explain Ark 564. For instance, consider a sparse, uniform distribution of spheres such that 10% of the sky as seen by the central X-ray source is covered by spheres. Neglecting second-order reflection, the spectrum of the reflected X-rays seen from a single direction would be as shown in Fig. 37 by the green spectrum. This spectrum is reminiscent of the ‘delayed’ spectrum of Section 4.1.




Fig. 37.— Reflected spectra from a uniform sphere of gas with maximum optical depth $\tau_c = 8\tau_0$ and temperature $kT = 100$ keV. The input spectrum (blue) is the XSPEC model `nthcomp` with a power-law of $\Gamma = 3.0$ and low and high energy cutoffs of 0.1 keV and 300 keV respectively. Output spectra for each of 100 evenly spaced bins in $\cos\theta$ are shown in shades of grey, with light-grey representing the forward directions and black the backward. A spectrum representative of a source covering fraction of 10% is shown in green.

It is notable that such a spectrum can be produced with just a 10% covering fraction. This is due to photons not being lost, as would be the case with absorption, but instead shifted to higher energies. Of course, the scenario envisaged is a gross simplification of what is likely to be the case in reality. Rather than neat, orderly, spheres the true situation likely involves irregular clouds, non-uniform

densities, and temperature gradients. The fast variability of the X-rays observed argues in favour of the covering gas being located in a small region around the central engine, with a length scale on the order of a thousand light-seconds (Boller et al. 1996). That such a simplistic model can approximate the observed behaviour is encouraging. Unfortunately, it is not practical to further home in on a more exact gas distribution. Firstly the data is of poor quality. Secondly, the Monte-Carlo approach is not suited to fitting since it requires a lengthy simulation for each set of trial parameters.

6. Conclusions and Avenues for Further Research

In Section 2, a novel algorithm was developed for estimating the response of X-ray time series to flares in activity. The method identifies flares in a single energy band, and finds the response to those flares in higher energy bands. The method is capable of coping with the gappy and noisy nature of X-ray time series, and stands up to testing against synthetic data.

In Section 3, the algorithm was applied to data from three AGN. In each case, flares in emission in soft energy bands are accompanied by corresponding flares in harder energy bands. In each case the response in the harder energy band shows signs of an excess in emission at later times than the flare, with Ark 564 showing the most prominent and dramatic excess. This delayed excess is used to conclude that the most likely explanation for the behaviour of Ark 564 is that of X-ray reverberation.

Section 4 further explores the nature of Ark 564, delving into a joint temporal and spectral analysis. Spectral components associated with the initial flare and the delayed excess were extracted. The extracted components were used in a simple top-hat reverberation model, and compared against a joint temporal and spectral evaluation of the lag-spectrum, with favourable results. Having concluded that the spectral components fit within a physically plausible and meaningful reverberation scenario, attention turned to finding an explanation for the spectra. Attempts

at fitting absorption models to the spectra failed to find physically convincing solutions. Instead, a model was tried in which the primary source spectrum is Compton upscattered by distant surrounding material. This model was found to be both physically plausible and easy to fit to the spectra.

In Section 5, a thorough simulation of reflection by Comptonizing regions was carried out. The evaluation of Comptonizing models in Section 4 had found a paucity of literature on the distribution of photons after reflection by Comptonizing regions. The Monte-Carlo simulations revealed a rich and complex variation of the reflected flux and spectral shape over varying incidence and reflection angles.

The overall conclusion of this work is that the behaviour of Ark 564 favours a model (illustrated in Fig. 38) in which the primary flaring source undergoes reverberation by hot Comptonizing material. The Comptonizing material has a covering fraction of roughly 10% and exists in a region roughly 1500 light-seconds from the source.

6.1. Further Research

The research performed here suggests several avenues for future work. First and foremost is the hope that another AGN with the properties of Ark 564 can be found. Such an object might be found among other ‘simple’ NLS1, characterised by Gallo (2006) as being those in which the X-ray spectrum can be well described with the most basic of spectral models. This is a classification which distinguishes Ark 564 from 1H 0707–496 and NGC 4051. Of particular interest might be the




Fig. 38.— A diagram showing the components of the preferred model for the components of Ark 564. a) The black hole. b) The accretion disk. c) The 'corona' X-ray source, 500 light-second lengthscale. d) Reflecting clouds of hot gas, 1500 light-seconds from the source, 10% covering fraction.

object PG 1244+026, examined by Jin et al. (2013), which has a smooth spectrum that can be modeled by components strikingly reminiscent of those found for Ark 564. The discovery of another object with such a clear delayed component as Ark 564 would provide valuable reassurance as to the validity of the results and conclusions reached. Second, a lengthy follow-up observation of Ark 564 would provide welcome clarification of several points. For instance the evidence for the existence of different classes of flares in emission could be firmed up, and

an explanation for the phenomenon reached. An additional observation would also allow for evaluation of the degree of non-stationarity present in the system, and discover whether the components identified in this work are persistent features.

In light of the discrepancy with respect to the results of Kara et al. (2013) noted in Section 4.2, it may be of interest to re-examine the published lag-energy spectra of other AGN than Ark 564. Of particular note might be the published lag-energy spectra of 1H 0707–495 due to that object’s prominence in the literature. Other objects for which lag-energy spectra were not calculated using the more rigorous maximum-likelihood method would also be of interest.

The work performed in Section 5 provides ample scope for further research. The simulations presented here have been restricted to two specific geometries (slabs and spheres), both of which are of dubious physical likelihood. There is a great deal of potential for testing different input spectra, more complex and realistic geometries, and variations in density and temperature across the distribution. It also has a broader relevance than the fairly narrow scope of X-ray reverberation, allowing a more generalised examination of Comptonization than that present in the current literature. There is also a need for research to identify the cause of the apparent discrepancy between the behaviour of the `comptt` and `nthcomp` models for Comptonization.

REFERENCES

- M. A. Abramowicz, M. Jaroszyński, S. Kato, J.-P. Lasota, A. Różańska, and A. Sądowski. Leaving the innermost stable circular orbit: the inner edge of a black-hole accretion disk at various luminosities. *A&A*, 521:A15, October 2010. doi: 10.1051/0004-6361/201014467.
- W. N. Alston, S. Vaughan, and P. Uttley. The flux-dependent X-ray time lags in NGC 4051. *MNRAS*, 435:1511–1519, October 2013. doi: 10.1093/mnras/stt1391.
- R. Antonucci. Unified models for active galactic nuclei and quasars. *ARA&A*, 31: 473–521, 1993. doi: 10.1146/annurev.aa.31.090193.002353.
- P. Arévalo and P. Uttley. Investigating a fluctuating-accretion model for the spectral-timing properties of accreting black hole systems. *MNRAS*, 367: 801–814, April 2006. doi: 10.1111/j.1365-2966.2006.09989.x.
- K. A. Arnaud. XSPEC: The First Ten Years. In G. H. Jacoby and J. Barnes, editors, *Astronomical Data Analysis Software and Systems V*, volume 101 of *Astronomical Society of the Pacific Conference Series*, page 17, 1996.
- T. Boller, W. N. Brandt, and H. Fink. Soft X-ray properties of narrow-line Seyfert 1 galaxies. *A&A*, 305:53, January 1996.
- T. A. Boroson. Black Hole Mass and Eddington Ratio as Drivers for the Observable Properties of Radio-loud and Radio-quiet QSOs. *ApJ*, 565:78–85, January 2002. doi: 10.1086/324486.

- V. Botte, S. Ciroi, P. Rafanelli, and F. Di Mille. Exploring Narrow-Line Seyfert 1 Galaxies through the Physical Properties of Their Hosts. *AJ*, 127:3168–3179, June 2004. doi: 10.1086/420803.
- W. Brinkmann, J. Siebert, W. Reich, E. Fuerst, P. Reich, W. Voges, J. Truemper, and R. Wielebinski. The ROSAT AGN content of the 87GB 5 GHz survey: bulk properties of previously optically identified sources. *A&AS*, 109: 147–170, January 1995.
- B. De Marco, G. Ponti, M. Cappi, M. Dadina, P. Uttley, E. M. Cackett, A. C. Fabian, and G. Miniutti. Discovery of a relation between black hole mass and soft X-ray time lags in active galactic nuclei. *MNRAS*, 431:2441–2452, May 2013. doi: 10.1093/mnras/stt339.
- K. D. Denney, L. C. Watson, B. M. Peterson, R. W. Pogge, D. W. Atlee, M. C. Bentz, J. C. Bird, D. J. Brokofsky, M. L. Comins, M. Dietrich, V. T. Doroshenko, J. D. Eastman, Y. S. Efimov, C. M. Gaskell, C. H. Hedrick, S. A. Klimanov, E. S. Klimek, A. K. Kruse, J. B. Lamb, K. Leighly, T. Minezaki, S. V. Nazarov, E. A. Petersen, P. Peterson, S. Poindexter, Y. Schlesinger, K. J. Sakata, S. G. Sergeev, J. J. Tobin, C. Unterborn, M. Vestergaard, A. E. Watkins, and Y. Yoshii. A Revised Broad-line Region Radius and Black Hole Mass for the Narrow-line Seyfert 1 NGC 4051. *ApJ*, 702:1353–1366, September 2009. doi: 10.1088/0004-637X/702/2/1353.
- C. Done, S. W. Davis, C. Jin, O. Blaes, and M. Ward. Intrinsic disc emission and the soft X-ray excess in active galactic nuclei. *MNRAS*, 420:1848–1860, March 2012. doi: 10.1111/j.1365-2966.2011.19779.x.

- B. Efron. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1):1–26, 01 1979. doi: 10.1214/aos/1176344552. URL <http://dx.doi.org/10.1214/aos/1176344552>.
- M. Elvis. A Structure for Quasars. *ApJ*, 545:63–76, December 2000. doi: 10.1086/317778.
- A. C. Fabian, A. Zoghbi, R. R. Ross, P. Uttley, L. C. Gallo, W. N. Brandt, A. J. Blustin, T. Boller, M. D. Caballero-Garcia, J. Larsson, J. M. Miller, G. Miniutti, G. Ponti, R. C. Reis, C. S. Reynolds, Y. Tanaka, and A. J. Young. Broad line emission from iron K- and L-shell transitions in the active galaxy 1H0707-495. *Nature*, 459:540–542, May 2009. doi: 10.1038/nature08007.
- A. C. Fabian, A. Zoghbi, D. Wilkins, T. Dwelly, P. Uttley, N. Schartel, G. Miniutti, L. Gallo, D. Grupe, S. Komossa, and M. Santos-Lleó. 1H 0707-495 in 2011: an X-ray source within a gravitational radius of the event horizon. *MNRAS*, 419:116–123, January 2012. doi: 10.1111/j.1365-2966.2011.19676.x.
- A. C. Fabian, M. L. Parker, D. R. Wilkins, J. M. Miller, E. Kara, C. S. Reynolds, and T. Dauser. On the determination of the spin and disc truncation of accreting black holes using X-ray reflection. *MNRAS*, 439:2307–2313, April 2014. doi: 10.1093/mnras/stu045.
- C.-A. Faucher-Giguère and E. Quataert. The physics of galactic winds driven by active galactic nuclei. *MNRAS*, 425:605–622, September 2012. doi: 10.1111/j.1365-2966.2012.21512.x.

- L. Ferrarese and H. Ford. Supermassive Black Holes in Galactic Nuclei: Past, Present and Future Research. *Space Sci. Rev.*, 116:523–624, February 2005. doi: 10.1007/s11214-005-3947-6.
- L. C. Gallo. Investigating the nature of narrow-line Seyfert 1 galaxies with high-energy spectral complexity. *MNRAS*, 368:479–486, May 2006. doi: 10.1111/j.1365-2966.2006.10137.x.
- L. C. Gallo, T. Boller, Y. Tanaka, A. C. Fabian, W. N. Brandt, W. F. Welsh, N. Anabuki, and Y. Haba. The X-ray variability of the narrow-line type 1 Seyfert galaxy IRAS 13224-3809 from an XMM-Newton observation. *MNRAS*, 347:269–276, January 2004. doi: 10.1111/j.1365-2966.2004.07196.x.
- E. Gardner and C. Done. A physical model for the X-ray time lags of narrow-line Seyfert type 1 active galactic nuclei. *MNRAS*, 442:2456–2473, August 2014. doi: 10.1093/mnras/stu1026.
- G. Ghisellini, L. Foschini, M. Volonteri, G. Ghirlanda, F. Haardt, D. Burlon, and F. Tavecchio. The blazar S5 0014+813: a real or apparent monster? *MNRAS*, 399:L24–L28, October 2009. doi: 10.1111/j.1745-3933.2009.00716.x.
- M. Giustini, T. J. Turner, J. N. Reeves, L. Miller, E. Legg, S. B. Kraemer, and I. M. George. Revealing a hard X-ray spectral component that reverberates within one light hour of the central supermassive black hole in Ark 564. *ArXiv e-prints*, February 2015.
- A. Gorecki and W. Wilczewski. A study of Comptonization of radiation in an

- electron plasma using the Monte Carlo method. *Acta Astron.*, 34:141–160, 1984.
- A. Gupta, S. Mathur, Y. Krongold, and F. Nicastro. Discovery of Relativistic Outflow in the Seyfert Galaxy Ark 564. *ApJ*, 772:66, July 2013. doi: 10.1088/0004-637X/772/1/66.
- F. Haardt and L. Maraschi. A two-phase model for the X-ray emission from Seyfert galaxies. *ApJ*, 380:L51–L54, October 1991. doi: 10.1086/186171.
- F. J. Harris. On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform. *IEEE Proceedings*, 66:51–83, 1978.
- A. Hewitt and G. Burbidge. An optical catalog of extragalactic emission-line objects similar to quasi-stellar objects. *ApJS*, 75:297–356, February 1991. doi: 10.1086/191533.
- X.-M. Hua and L. Titarchuk. Comptonization Models and Spectroscopy of X-Ray and Gamma-Ray Sources: A Combined Study by Monte Carlo and Analytical Methods. *ApJ*, 449:188, August 1995. doi: 10.1086/176045.
- C. Jin, C. Done, M. Middleton, and M. Ward. A long XMM-Newton observation of an extreme narrow-line Seyfert 1: PG 1244+026. *MNRAS*, 436:3173–3185, December 2013. doi: 10.1093/mnras/stt1801.
- T. Kallman. Photoionization Modeling using Xstar. In *American Astronomical Society Meeting Abstracts #206*, volume 37 of *Bulletin of the American Astronomical Society*, page 480, May 2005.

- E. Kara, A. C. Fabian, E. M. Cackett, P. Uttley, D. R. Wilkins, and A. Zoghbi.
Discovery of high-frequency iron K lags in Ark 564 and Mrk 335. *MNRAS*,
434:1129–1137, September 2013. doi: 10.1093/mnras/stt1055.
- S. Kaspi, P. S. Smith, H. Netzer, D. Maoz, B. T. Jannuzi, and U. Giveon.
Reverberation Measurements for 17 Quasars and the Size-Mass-Luminosity
Relations in Active Galactic Nuclei. *ApJ*, 533:631–649, April 2000. doi:
10.1086/308704.
- O. Klein and T. Nishina. Über die Streuung von Strahlung durch freie Elektronen
nach der neuen relativistischen Quantendynamik von Dirac. *Zeitschrift für
Physik*, 52:853–868, November 1929. doi: 10.1007/BF01366453.
- C. Knigge, J. A. Woods, and J. E. Drew. The application of Monte Carlo methods
to the synthesis of spectral line profiles arising from accretion disc winds.
MNRAS, 273:225–248, March 1995.
- J. Kormendy and D. Richstone. Inward Bound—The Search For Super-
massive Black Holes In Galactic Nuclei. *ARA&A*, 33:581, 1995. doi:
10.1146/annurev.aa.33.090195.003053.
- O. Kotov, E. Churazov, and M. Gilfanov. On the X-ray time-lags in the
black hole candidates. *MNRAS*, 327:799–807, November 2001. doi:
10.1046/j.1365-8711.2001.04769.x.
- S. Krucker and R. P. Lin. Relative Timing and Spectra of Solar Flare
Hard X-ray Sources. *Sol. Phys.*, 210:229–243, November 2002. doi:
10.1023/A:1022469902940.

- D. Kundur and D. Hatzinakos. Blind image deconvolution. *IEEE Signal Processing Magazine*, 13:43–64, May 1996. doi: 10.1109/79.489268.
- E. Legg, L. Miller, T. J. Turner, M. Giustini, J. N. Reeves, and S. B. Kraemer. Direct Measurement of the X-Ray Time-delay Transfer Function in Active Galactic Nuclei. *ApJ*, 760:73, November 2012. doi: 10.1088/0004-637X/760/1/73.
- K. M. Leighly. Hubble Space Telescope STIS Ultraviolet Spectral Evidence of Outflow in Extreme Narrow-Line Seyfert 1 Galaxies. II. Modeling and Interpretation. *ApJ*, 611:125–152, August 2004. doi: 10.1086/422089.
- K. M. Leighly and J. R. Moore. Hubble Space Telescope STIS Ultraviolet Spectral Evidence of Outflow in Extreme Narrow-Line Seyfert 1 Galaxies. I. Data and Analysis. *ApJ*, 611:107–124, August 2004. doi: 10.1086/422088.
- N. R. Lomb. Least-squares frequency analysis of unequally spaced data. *Ap&SS*, 39:447–462, February 1976. doi: 10.1007/BF00648343.
- M. S. Longair. *High energy astrophysics. Vol.1: Particles, photons and their detection*. March 1992.
- Y. E. Lyubarskii. Flicker noise in accretion discs. *MNRAS*, 292:679, December 1997.
- P. Magdziarz and A. A. Zdziarski. Angle-dependent Compton reflection of X-rays and gamma-rays. *MNRAS*, 273:837–848, April 1995.
- J. Magorrian, S. Tremaine, D. Richstone, R. Bender, G. Bower, A. Dressler, S. M. Faber, K. Gebhardt, R. Green, C. Grillmair, J. Kormendy, and T. Lauer.

- The Demography of Massive Dark Objects in Galaxy Centers. *AJ*, 115: 2285–2305, June 1998. doi: 10.1086/300353.
- R. Maiolino and G. H. Rieke. Low-Luminosity and Obscured Seyfert Nuclei in Nearby Galaxies. *ApJ*, 454:95, November 1995. doi: 10.1086/176468.
- I. McHardy. X-Ray Variability of AGN and Relationship to Galactic Black Hole Binary Systems. In T. Belloni, editor, *Lecture Notes in Physics, Berlin Springer Verlag*, volume 794 of *Lecture Notes in Physics, Berlin Springer Verlag*, page 203, March 2010. doi: 10.1007/978-3-540-76937-8_8.
- I. M. McHardy, P. Arévalo, P. Uttley, I. E. Papadakis, D. P. Summons, W. Brinkmann, and M. J. Page. Discovery of multiple Lorentzian components in the X-ray timing properties of the Narrow Line Seyfert 1 Ark 564. *MNRAS*, 382:985–994, December 2007. doi: 10.1111/j.1365-2966.2007.12411.x.
- D. B. Melrose, editor. *Quantum Plasmadynamics*, volume 735 of *Lecture Notes in Physics, Berlin Springer Verlag*, 2008. doi: 10.1007/978-0-387-73903-8.
- L. Miller and T. J. Turner. X-ray reverberation in NLS1. In *Narrow-Line Seyfert 1 Galaxies and their Place in the Universe*, 2011.
- L. Miller, T. J. Turner, J. N. Reeves, I. M. George, S. B. Kraemer, and B. Wingert. The variable X-ray spectrum of Markarian 766. I. Principal components analysis. *A&A*, 463:131–143, February 2007. doi: 10.1051/0004-6361:20066548.

- L. Miller, T. J. Turner, J. N. Reeves, and V. Braito. X-ray reverberation in 1H0707-495 revisited. *MNRAS*, 408:1928–1935, November 2010a. doi: 10.1111/j.1365-2966.2010.17261.x.
- L. Miller, T. J. Turner, J. N. Reeves, A. Lobban, S. B. Kraemer, and D. M. Crenshaw. Spectral variability and reverberation time delays in the Suzaku X-ray spectrum of NGC 4051. *MNRAS*, 403:196–210, March 2010b. doi: 10.1111/j.1365-2966.2009.16149.x.
- G. Miniutti and A. C. Fabian. A light bending model for the X-ray temporal and spectral properties of accreting black holes. *MNRAS*, 349:1435–1448, April 2004. doi: 10.1111/j.1365-2966.2004.07611.x.
- S. Miyamoto, S. Kitamoto, K. Mitsuda, and T. Dotani. Delayed hard X-rays from Cygnus X-1. *Nature*, 336:450–452, December 1988. doi: 10.1038/336450a0.
- R. Morrison and D. McCammon. Interstellar photoelectric absorption cross sections, 0.03-10 keV. *ApJ*, 270:119–122, July 1983. doi: 10.1086/161102.
- K. Nandra, I. M. George, R. F. Mushotzky, T. J. Turner, and T. Yaqoob. ASCA Observations of Seyfert 1 Galaxies. I. Data Analysis, Imaging, and Timing. *ApJ*, 476:70–82, February 1997.
- M. A. Nowak, B. A. Vaughan, J. Wilms, J. B. Dove, and M. C. Begelman. Rossi X-Ray Timing Explorer Observation of Cygnus X-1. II. Timing Analysis. *ApJ*, 510:874–891, January 1999. doi: 10.1086/306610.
- D. E. Osterbrock and R. W. Pogge. The spectra of narrow-line Seyfert 1 galaxies. *ApJ*, 297:166–176, October 1985. doi: 10.1086/163513.

- I. E. Papadakis, K. Nandra, and D. Kazanas. Frequency-dependent Time Lags in the X-Ray Emission of the Seyfert Galaxy NGC 7469. *ApJ*, 554:L133–L137, June 2001. doi: 10.1086/321722.
- B. M. Peterson. Reverberation mapping of active galactic nuclei. *PASP*, 105: 247–268, March 1993.
- L. A. Pozdniakov, I. M. Sobol, and R. A. Siuniae. Effect of multiple Compton scattering on an X-ray spectrum - Monte Carlo calculations. *AZh*, 54: 1246–1258, December 1977.
- D. Proga. Theory of Winds in AGNs. In L. C. Ho and J.-W. Wang, editors, *The Central Engine of Active Galactic Nuclei*, volume 373 of *Astronomical Society of the Pacific Conference Series*, page 267, October 2007.
- D. Proga, J. M. Stone, and T. R. Kallman. Dynamics of Line-driven Disk Winds in Active Galactic Nuclei. *ApJ*, 543:686–696, November 2000. doi: 10.1086/317154.
- M. J. Rees. Appearance of Relativistically Expanding Radio Sources. *Nature*, 211: 468–470, July 1966. doi: 10.1038/211468a0.
- J. D. Scargle. Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data. *ApJ*, 263:835–853, December 1982. doi: 10.1086/160554.
- J. D. Scargle. Studies in astronomical time series analysis. III - Fourier transforms, autocorrelation functions, and cross-correlation functions of unevenly spaced data. *ApJ*, 343:874–887, August 1989. doi: 10.1086/167757.

- N. I. Shakura and R. A. Syunyaev. Black holes in binary systems. Observational appearance. *A&A*, 24:337–355, 1973.
- S. A. Sim, K. S. Long, L. Miller, and T. J. Turner. Multidimensional modelling of X-ray spectra for AGN accretion disc outflows. *MNRAS*, 388:611–624, August 2008. doi: 10.1111/j.1365-2966.2008.13466.x.
- J. F. Steiner, R. Narayan, J. E. McClintock, and K. Ebisawa. A Simple Comptonization Model. *PASP*, 121:1279–1290, November 2009. doi: 10.1086/648535.
- R. A. Sunyaev and L. G. Titarchuk. Comptonization of X-rays in plasma clouds - Typical radiation spectra. *A&A*, 86:121–138, June 1980.
- R. A. Sunyaev and L. G. Titarchuk. Comptonization of low-frequency radiation in accretion disks Angular distribution and polarization of hard radiation. *A&A*, 143:374–388, February 1985.
- A. Takahata, E. Nadalin, R. Ferrari, L. Duarte, R. Suyama, R. Lopes, J. Romano, and M. Tygel. Unsupervised Processing of Geophysical Signals: A Review of Some Key Aspects of Blind Deconvolution and Blind Source Separation. *IEEE Signal Processing Magazine*, 29:27–35, July 2012. doi: 10.1109/MSP.2012.2189999.
- J. Terrell. Quasi-Stellar Diameters and Intensity Fluctuations. *Science*, 145: 918–919, August 1964. doi: 10.1126/science.145.3635.918.
- L. Titarchuk. Generalized Comptonization models and application to the

- recent high-energy observations. *ApJ*, 434:570–586, October 1994. doi: 10.1086/174760.
- L. G. Titarchuk and X.-M. Hua. Comptonization models and spectroscopy of X-ray and gamma-ray sources. *Advances in Space Research*, 19:99–108, May 1997. doi: 10.1016/S0273-1177(97)00043-4.
- F. Tombesi, M. Cappi, J. N. Reeves, G. G. C. Palumbo, T. Yaqoob, V. Braito, and M. Dadina. Evidence for ultra-fast outflows in radio-quiet AGNs. I. Detection and statistical incidence of Fe K-shell absorption lines. *A&A*, 521:A57, October 2010. doi: 10.1051/0004-6361/200913440.
- F. Tombesi, M. Cappi, J. N. Reeves, R. S. Nemmen, V. Braito, M. Gaspari, and C. S. Reynolds. Unification of X-ray winds in Seyfert galaxies: from ultra-fast outflows to warm absorbers. *MNRAS*, 430:1102–1117, April 2013. doi: 10.1093/mnras/sts692.
- T. J. Turner and L. Miller. X-ray absorption and reflection in active galactic nuclei. *A&A Rev.*, 17:47–104, March 2009. doi: 10.1007/s00159-009-0017-1.
- T. J. Turner, I. M. George, K. Nandra, and D. Turcan. On X-Ray Variability in Seyfert Galaxies. *ApJ*, 524:667–673, October 1999. doi: 10.1086/307834.
- M. van der Klis, G. Hasinger, L. Stella, A. Langmeier, J. van Paradijs, and W. H. G. Lewin. The complex cross-spectra of Cygnus X-2 and GX 5-1. *ApJ*, 319:L13–L18, August 1987. doi: 10.1086/184946.
- J. Wilms, A. Allen, and R. McCray. On the Absorption of X-Rays in the Interstellar Medium. *ApJ*, 542:914–924, October 2000. doi: 10.1086/317016.

- A. A. Zdziarski, W. N. Johnson, and P. Magdziarz. Broad-band γ -ray and X-ray spectra of NGC 4151 and their implications for physical processes and geometry. *MNRAS*, 283:193–206, November 1996.
- X.-L. Zhou and J.-M. Wang. Narrow Iron $K\alpha$ Lines in Active Galactic Nuclei: Evolving Populations? *ApJ*, 618:L83–L86, January 2005. doi: 10.1086/427871.
- A. Zoghbi, A. C. Fabian, P. Uttley, G. Miniutti, L. C. Gallo, C. S. Reynolds, J. M. Miller, and G. Ponti. Broad iron L line and X-ray reverberation in 1H0707-495. *MNRAS*, 401:2419–2432, February 2010. doi: 10.1111/j.1365-2966.2009.15816.x.
- A. Zoghbi, P. Uttley, and A. C. Fabian. Understanding reverberation lags in 1H0707-495. *MNRAS*, 412:59–64, March 2011. doi: 10.1111/j.1365-2966.2010.17883.x.
- A. Zoghbi, C. Reynolds, and E. M. Cackett. Calculating Time Lags from Unevenly Sampled Light Curves. *ApJ*, 777:24, November 2013. doi: 10.1088/0004-637X/777/1/24.