X-Ray Power Density Spectra of Black Hole Binaries: A New Deadtime Model for the RXTE PCA

by

Dennis Wei

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Bachelor of Science in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2006

Author . . . .
Department of Physics
May 12, 2006

Certified by. . . . .
Ronald A. Remillard
Principal Research Scientist
Thesis Supervisor

Accepted by . . . .
David E. Pritchard
Senior Thesis Coordinator, Department of Physics

X-Ray Power Density Spectra of Black Hole Binaries: A New Deadtime Model for the RXTE PCA

by

Dennis Wei

Abstract

The power density spectrum is an essential tool for determining the frequency content of X-ray radiation from astronomical sources. For neutron star systems, power density spectra reveal coherent oscillations for those sources that are pulsars, while quasi-periodic oscillations over a wide range of frequencies (0.01 to 1300 Hz) are used to identify subclasses and to probe the details of accretion physics. For black hole binaries, the power density spectrum is useful in many important contexts: distinguishing black hole binaries from neutron star binaries, tracking the evolution of X-ray states, and understanding the dynamics of accretion disks, in particular the high-frequency oscillations that appear to be rooted in general relativity for strong gravitational fields.

However, measurements of the power density spectrum are modified by the effects of deadtime in X-ray detectors. In this work, we focus on the Proportional Counter Array (PCA) instrument of the Rossi X-ray Timing Explorer (RXTE), an orbiting observatory that offers fast, microsecond-level time resolution and modest spectral resolution for celestial X-ray sources. We derive a new model for the effect of detector deadtime on measurements of the power density spectrum. The model treats in a unified manner the contributions from self-deadtime among selected events and interference from non-selected events. Using high-frequency power density spectra obtained from observations of X-ray sources, the new model is shown to be more accurate than existing approaches. The comparison between the model and the observations leads to a measurement of 8.83 μs for the fundamental instrument deadtime timescale, which is dominated by the analog-to-digital conversion time. We additionally measure 59 μs and 137 μs for the Very Large Event deadtime related to observer-specified settings 1 and 2 respectively. Future refinements to the deadtime model are discussed, such as corrections for highly variable sources and for individual X-ray energy bands.

A preliminary comparison between power density spectra from black hole binaries and neutron star binaries is undertaken using the new deadtime model. While it may be possible to use high-frequency cut-offs in the power continuum to distinguish neu-

tron star binaries from black hole binaries in the thermal and hard X-ray states, the comparison is inconclusive for black hole binaries in the steep power-law state. Since state definitions require considerations of X-ray spectral properties, the comparison results dispute a suggestion in the literature that accreting neutron stars and black holes can be distinguished on the basis of power density spectra alone.

Thesis Supervisor: Ronald A. Remillard
Title: Principal Research Scientist

Acknowledgments

I would like to thank my thesis advisor, Ron Remillard, for taking me on as a senior thesis student, for his enthusiasm for the project, and for his total commitment to me. We have had many deep discussions that I found greatly enlightening. I especially value the confidence he placed in me to pursue my own ideas and the collegial respect he showed for my opinions. In addition, Ron was patient and painstakingly diligent in reading and helping to revise this thesis, dramatically improving its scope, clarity, and overall quality.

I would also like to thank Ed Morgan, research scientist at the MIT Kavli Institute (MKI) for Astrophysics and Space Research. Ed was always available to answer my questions and his contributions during group discussions were insightful. I don't know how he managed to fix our software problems so quickly and efficiently. Thanks also to Dacheng Lin, graduate student at MKI, for his thorough revisions to our spectral analysis software, and to Brian Pepper, MIT undergraduate, for his care in helping with the data analysis.

Lastly, to my parents, for their unwavering love and support. They have always, through gentle encouragement, led me on a path toward higher goals and greater success. During my time at MIT, they have taken care of or reminded me of many of my obligations so that I could concentrate on my studies. Their wisdom and judgement were indispensable for those larger, weightier problems outside of academics, the ones that cannot be solved through equations. I cannot hope to repay or even adequately express all that I owe to them.

Contents

List of Figures

10

List of Tables

Chapter 1

Introduction and Background

1.1   Power density spectra in X-ray astronomy

In X-ray astronomy, the power density spectrum (PDS) is an indispensable tool for describing and analyzing the frequency content of X-ray radiation. Many celestial X-ray sources are compact objects that accrete matter from surrounding structures such as companion stars in a close binary system (see Psaltis (2006) for background). The accretion process drives the emission of X-rays, which we observe using X-ray detectors. The X-ray flux is often irregular, turbulent, and contains signatures of many different timescales, a reflection of the complex dynamics of accretion (e.g. Mc-Clintock & Remillard 2006). Consequently, observations of X-ray sources in the form of light curves, which record the number of photon arrivals as a function of time, are difficult to interpret directly. However, transforming light curves into the power spectral domain, which characterizes the flux as a function of frequency, affords a clearer understanding of the dynamical timescales and fast fluctuations of the accreting system. Power spectral measurements can therefore yield many insights into the behaviour of accreting compact objects (van der Klis 2006).

The astronomical phenomena that are analyzed using the PDS may be broken into two broad categories: periodic and aperiodic variability. Examples of periodic variability include the well-known pulsations of strongly magnetized, rotating neutron stars known as pulsars. Periodic oscillations are more easily identified and charac-

terized in the power spectral domain, especially in cases where the oscillation has a low amplitude or a complicated profile. Flux oscillations are often produced by binary system dynamics, e.g. the offset magnetic pole of a pulsar spinning about its rotational axis. With the aid of the PDS, oscillation frequencies, amplitudes, harmonic relationships and other features may be measured and information about the responsible mechanisms may be gained.

Aperiodic variability encompasses any fluctuations that do not repeat at regular intervals. Aperiodic variability is more common than periodic variability because accretion flows are generally turbulent as already mentioned. Characterizing the aperiodic variability of an X-ray source is an important step in understanding the often complicated dynamics of the accretion disks that surround compact objects. The study of aperiodic variability is often concerned with both the shape and the integrated amplitude of the continuum PDS (Remillard & McClintock 2006). It is common to describe the shape of the continuum in different frequency bands in terms of power-law or broken power-law relationships. Alternatively, the continuum may be modelled as the superposition of several very broad Lorentzian functions. The integrated amplitude in a frequency band measures the flux level in the corresponding range of timescales. In addition, the presence of a cut-off frequency, beyond which no meaningful power can be found, is an important indicator of the dynamical limits of the radiating process. Cut-off frequencies are discussed as a means of distinguishing black holes from neutron stars in Section 1.3.

Aperiodic variability includes quasi-periodic oscillations (QPOs), which are peaks of power concentrated in narrow frequency ranges but are less coherent than periodic oscillations (van der Klis 2006). Analysis of the PDS allows the centre frequency, integrated power, and coherence of a QPO to be measured. QPOs are often interpreted as signatures of resonances within accretion disks; their significance for black hole binaries will be discussed in Section 1.2.

PDSs derived from real observations are composed of two components: a component due to actual variations in source flux, and a component due to Poisson counting statistics. We will also refer to the source component as the source power or

16

the source PDS, and to the Poisson statistical component as the Poisson power, the Poisson power level, or the Poisson PDS. The Poisson component is a consequence of the finite collecting area of real instruments and the discrete nature of X-ray measurements (e.g. counts per time interval). The source and Poisson components will be discussed in greater detail in Section 4.1. X-ray astronomers are interested in variations in source flux and not in Poisson statistics, so for scientific studies the source power must be isolated. In order to extract the source power, it is necessary to model the Poisson power and subtract it from the total PDS. The quality of the model for the Poisson power therefore limits the accuracy to which the source power may be determined.

The aim of this thesis is to provide a more accurate model of the Poisson statistical contribution to the total PDS in the presence of detector deadtime effects. Detector deadtime refers to the time required by the signal processing elements in X-ray detectors to recover from an incident event. During this processing time, no new events are recorded by the detector and the detector is considered "dead". Detector deadtime alters the Poisson PDS in a non-trivial, frequency-dependent manner. A precise and detailed model is therefore needed to accurately capture the deadtime-modified Poisson power level and remove it from the total PDS.

The model derived in this thesis is specific to the Proportional Counter Array (PCA) aboard the Rossi X-ray Timing Explorer (RXTE), a satellite-based observatory that offers fast, microsecond-level time resolution and modest spectral resolution for celestial X-ray sources. RXTE is currently the instrument of choice for fast X-ray timing studies because it has the largest collecting area (~ 6000 cm^2) of any X-ray instrument used for astrophysics (Jahoda et al. 2006; Swank 1998). It is hoped that the analytical techniques employed in this thesis will continue to be relevant for data collected by the next generation of observational instruments.

Many areas of astrophysical research rely upon power spectral analyses and could stand to benefit from an improved model for the Poisson component of the PDS. This thesis is not intended as a comprehensive description of PDS applications. Nevertheless, in Sections 1.2 and 1.3 we highlight two areas of recent activity where a more

careful calculation of PDSs could have a major impact.

1.2   Black hole binaries and X-ray states

Black hole binaries (BHBs) are important astronomical objects and are represented in many contemporary areas of inquiry. A black hole (BH), according to general relativity, has an event horizon bounding the region of spacetime that is unable to communicate with the external universe because of the strength of the BH's gravity (Shapiro & Teukolsky 1983). They are one possible outcome of the evolution and gravitational collapse of very massive stars (Woosley et al. 2002). A BHB consists of a BH and a companion star; the latter often contributes stellar material that is pulled by gravity into an accretion disk surrounding the BH. Currently, 20 X-ray binaries are confirmed to be BHBs, a small sample of a possible $\sim 10^9$ BHs present in our galaxy (Remillard & McClintock 2006).

Significant questions persist concerning the formation and evolution of BHBs and the accretion disk physics that underlies their behaviour. In addition, BHBs provide an excellent testing ground for general relativity in strong gravitational fields. Fortunately, with inner accretion disk temperatures reaching 1 keV (10^7 K), many BHBs are bright X-ray sources that can be studied using X-ray cameras on space missions such as RXTE. The timing and spectral data gathered from these missions play a major role in addressing the questions above (Remillard & McClintock 2006).

BHBs often exhibit spectral evolution that involves very different types of X-ray energy spectra and PDSs. The principal behaviour patterns are often classified in terms of X-ray "states". X-ray state definitions were traditionally based upon the source luminosity and the X-ray energy spectrum, which is usually composed of a thermal and a non-thermal component. The generally softer (lower X-ray energy) thermal component is modelled by a multi-temperature black body with a characteristic temperature near 1 keV (see description of thermal state later in this section). The harder (higher X-ray energy) non-thermal component is modelled by a power-law, $N(E) \propto E^{-\Gamma}$, where N is the number density of photons, E is the photon energy

and the exponent Γ is called the photon index.

In light of challenges to the traditional state definitions brought by RXTE observations, a new set of X-ray state definitions was developed (Remillard & McClintock 2006; McClintock & Remillard 2006). The new state definitions abandon source luminosity as a criterion and quantify state classifications in terms of parameters from both the X-ray energy spectrum and the PDS. Remillard & McClintock used the following four measures: the ratio f between the accretion disk flux and the total source flux from 2 to 20 keV, a measure of the proportion of thermal radiation; the photon index Γ; the variability r integrated from 0.1 Hz to 10 Hz in the PDS and expressed as a fraction of the mean source count rate; and the integrated amplitude a of any QPOs in the range $0.1 - 30$ Hz, also expressed as a fraction of the mean source count rate. Given these criteria, the three X-ray states for outbursting BHBs are as follows:

- *Thermal state*, defined by $f > 0.75$ (disk-dominated flux), $r < 0.075$ and $a_{max} < 0.005$ (low continuum power and weak QPOs).

- *Hard state*, defined by $f < 0.2$ (power-law-dominated flux), $1.4 < \Gamma < 2.1$, and $r > 0.1$ (high continuum power).

- *Steep power-law (SPL) state*, defined by $\Gamma > 2.4$, $r < 0.15$, and either $f < 0.8$ with a QPO present ($a > 0.01$) or $f < 0.5$ with no QPOs.

States that do not fall within the above ranges are classified as intermediate states.

Two of the four criteria defining X-ray states, the normalized source variability r and the QPO amplitude a, depend on the PDS, and specifically on the source component only. If the subtraction of the Poisson component is not performed accurately, the resulting estimates of r and a will be biased. In the context of X-ray states, it is therefore important to carefully model the effect of deadtime on Poisson power so that the underlying source power can be properly estimated.

The thermal state is dominated by thermal radiation from the inner regions of the accretion disk. The observed X-ray energy spectra in the thermal state are well-modelled by a multi-temperature disk radiating as a black body (Makishima et al.

1986; Kubota & Makishima 2004; Kubota et al. 2005). These models are based on the well-known steady-state, thin disk model of Shakura & Sunyaev (1973) for in-falling matter in an accretion disk around a compact object. The temperature profile of the disk varies inversely with radius; thus the innermost annuli of the disk are the strongest contributors to the radiation. PDSs drawn from the thermal state tend to decrease steadily with frequency in a manner resembling a power-law.

The hard state is strongly correlated with the presence of a steady radio jet from the same source (Fender 2006; McClintock & Remillard 2006). The evidence for the relationship includes VLBI images of radio jets in BHBs (Dhawan et al. 2000; Stirling et al. 2001) and the quenching of strong radio emission coinciding with transitions out of the hard X-ray state (Fender et al. 1999). PDSs of BHBs in the hard state tend to be elevated and roughly constant over a wide range of frequencies before dropping quickly after a maximum frequency around 10–100 Hz. The broad continuum in the X-ray PDS is thought to be caused by the same non-thermal mechanism that sustains the jet.

The SPL state is the least understood of the three X-ray states for BHBs. Most models suggest that inverse Compton scattering, the scattering of lower-energy photons by relativistic electrons, is the primary source of radiation (e.g. Zdziarski & Gierliński 2004). The seed photons originate from the accretion disk, and the scattering is thought to occur in the surrounding non-thermal corona.

One of the main signatures of the SPL state is the presence of QPOs, clearly apparent in the PDS as peaks distinct from the continuum. Low-frequency QPOs (centre frequency $\sim$ 0.1–30 Hz) can have high amplitudes ($a > 0.15$) and coherence values (e.g. Muno et al. 1999). With amplitudes that peak at photon energies above 10 keV, they are clearly correlated with the non-thermal component of the X-ray spectrum (Sobczak et al. 2000; Vignarca et al. 2003). There exist many models attributing low-frequency QPOs to different types of oscillations involving the accretion disk and/or magnetic instabilities (e.g. Titarchuk & Osherovich 2000; Chakrabarti & Manickam 2000; Nobili et al. 2000; Tagger & Pellat 1999).

On the other hand, high-frequency QPOs (centre frequencies $\sim$ 40–450 Hz) are

particularly interesting because their frequencies are comparable to the orbital frequency of the innermost stable circular orbit (ISCO) around a $\sim 10 M_\odot$ spinless BH. According to general relativity, the ISCO lies outside the event horizon and depends on the mass and the spin of the BH (Shapiro & Teukolsky 1983). The centre frequencies of high-frequency QPOs do not vary significantly with changes in luminosity, suggesting that they are intrinsic properties of the BH (Remillard et al. 2002, 2006). High-frequency QPOs could offer a method for measuring the spin of a BH when the mass is well-constrained from analysis of the binary motion of the companion star (Remillard & McClintock 2006). Of special interest is the discovery of high-frequency QPO pairs with frequencies in a 3 : 2 ratio, representing an exciting opportunity to test general relativistic theories of resonant oscillations (Abramowicz & Kluźniak 2001; Remillard & McClintock 2006).

Unfortunately, the identification and measurement of high-frequency QPOs in the hundreds of Hz range is made difficult by the intrinsically low amplitudes in that regime. An inaccurate subtraction of the Poisson power can severely distort the profiles of low-amplitude, high-frequency QPOs. In order to study high-frequency QPOs with greater confidence, one must first construct a trustworthy model for the Poisson power so that the inaccuracies in the model do not obscure high-frequency features in the source PDS.

1.3 Distinguishing neutron star binaries from black hole binaries

Further motivation for more carefully modelling the Poisson power comes from Sunyaev & Revnivtsev (2000), who have proposed a method for distinguishing neutron star binaries (NSBs) from BHBs by comparing their power spectral profiles. Sunyaev & Revnivtsev acknowledge that a number of criteria based on energy spectrum characteristics are already commonly employed. In addition, the detection of coherent X-ray pulsations is fairly definitive evidence for the offset rotating magnetic axis

of a pulsar, and the detection of Type I X-ray bursts, thermonuclear explosions of accreted hydrogen and/or helium, suggests the hard surface of a neutron star (NS). However, classification based on X-ray pulsations or bursts is not possible for sources that do not demonstrate such behaviour, i.e. not all accreting NSBs exhibit pulsations or bursts. Thus, classification using the PDS could provide a useful alternative to existing approaches. These comparisons between the PDSs of BHBs and NSBs may also provide indirect evidence for the existence of an ISCO for BHs.

Sunyaev & Revnivtsev compared PDSs obtained from 9 NSBs and 9 BHBs in the hard spectral state. The PDSs were normalized to units of squared fractional variability (RMS normalization, see Section 3.3) and then multiplied by the frequency, yielding plots of power $\times$ frequency versus frequency. Given this normalization, they observed that the NSB PDSs contain continuum power in the frequency range $500 - 1000$ Hz. In contrast, the BHB PDSs fall off rapidly for frequencies beyond 10 Hz, becoming one to two orders of magnitude lower than the NSB PDSs at frequencies beyond 100 Hz. From their empirical work, Sunyaev & Revnivtsev proposed identifying sources with significant power up to several hundred Hz as NSBs. At a very basic level, they attribute the difference in power spectral content to the different locations of principal energy release in NSBs and BHBs. Most of the radiation from a BHB is released in the accretion disk, while the boundary region between the inner edge of the accretion disk and the NS surface is the most active region in a NSB.

The method of Sunyaev & Revnivtsev (2000), or any other work that relies upon the high-frequency decay and cut-off of PDSs, is strongly sensitive to the model used to remove the deadtime-modified Poisson component of the PDS. At high frequencies, the intrinsic source power is almost always low, and becomes by definition even lower beyond a cut-off frequency. It is very difficult in this regime to separate the source power from the Poisson power. We indicate in Section 1.4 why the deadtime model used by Sunyaev & Revnivtsev is not very satisfying. As an alternative, we describe in Chapter 7 preliminary results from comparing BHB and NSB PDSs using the deadtime model developed in this thesis.

Furthermore, the selection of X-ray states for BHBs was incomplete in Sunyaev & Revnivtsev (2000), as they used observations of BHBs in the hard state only. A comparison of BHBs and NSBs should take into account the revised state definitions and extensive observational classifications that have been published in recent years (McClintock & Remillard 2006; Remillard & McClintock 2006). The dependence of high-frequency PDSs on the X-ray state is further addressed in Chapter 7.

1.4   Previous work in deadtime modelling

Detector deadtime effects were recognized and analyzed early on in the field of nuclear physics (see Evans (1955) for a textbook treatment). In X-ray astrophysics however, the deadtime problem did not surface until the EXOSAT mission, when it was taken up by a number of authors who derived approximations in the limits of low count rate or very short deadtime interval (e.g. Tennant 1987; Lewin et al. 1988; van der Klis 1989). A few years later, more rigorous analyses of the effect of deadtime on the PDS were contributed by Vikhlinin et al. (1994) and Zhang et al. (1995). This thesis builds mainly on the treatment of Zhang et al. (1995).

Zhang et al. distinguish between two classes of deadtime: paralyzable and non-paralyzable. In the case of paralyzable deadtime, an event that arrives while the detector is dead extends the deadtime for another interval from the moment it arrives. In theory, if the arrival rate is very high, the detector could be "paralyzed" indefinitely as events continue to engage it. In contrast, in the case of non-paralyzable deadtime, an event that arrives while the detector is dead has no effect at all. A single incidence of deadtime lasts for one interval and no longer.

Zhang et al. assume that the incident events arrive according to a Poisson process characterized by an incident rate r_{in} that cannot be observed directly. They relate r_{in} to the observed event rate r_o for both cases of paralyzable and non-paralyzable deadtime. They derive analytically the PDS predicted for a Poisson process undergoing deadtime of both varieties. The PDS for a Poisson process with paralyzable

deadtime interval t_d is given in terms of cyclic frequency f [Hz] as

$$P(f) = 2 - 4r_o \frac{\sin(2\pi f t_d)}{2\pi f}, \tag{1.1}$$

while the PDS for non-paralyzable deadtime is given by

$$P(f) = 2 - 4 \frac{r_{in}^2[1 - \cos(2\pi f t_d)] + 2\pi f r_{in} \sin(2\pi f t_d)}{r_{in}^2[1 - \cos(2\pi f t_d)]^2 + [r_{in}\sin(2\pi f t_d) + 2\pi f]^2}. \tag{1.2}$$

In the Leahy normalization (defined in Section 3.3), a pure Poisson process has an expected power level of 2; the terms following the 2 in equations (1.1) and (1.2) represent the deviation due to deadtime. Zhang et al. also derive expressions for the corresponding discrete-time PDS in the realistic case where events are recorded in terms of discrete time bins. Their calculations are verified by simulating an incident Poisson process subject to detector deadtime and comparing the simulated PDSs to their predictions.

The work of Zhang et al. (1995) and others describes the deadtime that the events of interest (termed "selected events") inflict upon themselves. There also exist other forms of deadtime caused by "non-selected events", i.e. events other than the selected ones, which the detector must also process. A special class of non-selected events, called very large events (VLEs), is particular to the RXTE PCA. VLEs consist of energetic particle events exceeding 120 keV in equivalent X-ray ionizing energy for the current PCA gain setting. To prevent VLEs from saturating the detector electronics and skewing the measurement of subsequent X-ray events, the detector is disabled following a VLE for an amount of time, τ_{VLE}, called the VLE window. The VLE window, which is a longer deadtime interval than the deadtime t_d for selected events, allows a full reset of the measurement chain. Further discussion can be found in Section 2.2. The deadtime effect of VLEs was first studied by Zhang et al. (1996). VLEs create anti-shots in the light curve and their effect on the Poisson PDS was modelled as an ad hoc additive term:

$$P(f) = 2r_0 r_{\mathrm{VLE}} \tau_{\mathrm{VLE}}^2 \left[\frac{\sin(\pi f \tau_{\mathrm{VLE}})}{\pi f \tau_{\mathrm{VLE}}} \right]^2, \qquad (1.3)$$

where r_{VLE} is the rate of VLEs per second.

Over the past decade, the combined deadtime models of Zhang et al. (1995) and Zhang et al. (1996) have been used extensively to correct PDSs obtained from RXTE PCA data (e.g. Morgan et al. 1997). Nevertheless, there are many cases where the Poisson power level is underestimated and the subtraction results in a positive residual that contaminates the high-frequency portions of the source PDS. A criticism of the combined Zhang et al. model is that it assumes the deadtime effects of selected events and VLEs to be additive. In reality, the two deadtime processes interact with each other and are best modelled together rather than separately. Moreover, the RXTE PCA event analyzers measure other types of non-selected events such as propane events, which are ionizations in a separate propane chamber sharing measurement circuitry with the main xenon chamber that generates selected events. Thus, non-selected events other than VLEs contribute to the deadtime experienced by selected events, but they have not received attention in existing deadtime models. The different categories of non-selected events will be described in greater detail in Section 2.2.

A different approach to correcting for deadtime effects has been put forth by Jernigan et al. (2000). Notably, Sunyaev & Revnivtsev (2000) employed a very similar method in their work on high-frequency PDSs. Jernigan et al. determine the deadtime correction by empirically fitting a model to PDSs at high frequencies, instead of predicting the correction based on observed count rates and deadtime parameters. In addition to a model that describes features of the source PDS (e.g. power-law slopes, QPO parameters), Jernigan et al. included two coefficients, A and B, that parameterize the deviation in the Poisson power caused by deadtime. The A coefficient controls the deadtime contribution from selected events, assumed to be paralyzable as in equation (1.1) without the leading 2. The B coefficient controls the VLE deadtime contribution as given by equation (1.3). The observed PDS is fit to a model that

accounts for both source and Poisson power components. The Poisson component is effectively removed by including it in the fit.

Revnivtsev et al. (2000) also adopted the method of fitting the observed PDS to a combined source and deadtime-Poisson model. However, unlike Jernigan et al. (2000), they assumed that the deadtime due to selected events is non-paralyzable and used equation (1.2) instead of equation (1.1).

The main drawback in the methods of Jernigan et al. (2000), Sunyaev & Revnivtsev (2000), and Revnivtsev et al. (2000) is the assumption that the source power disappears above a certain frequency, leaving only the Poisson power. While the assumption is generally true, it is possible that faint, high-frequency source power could be mistaken for Poisson power in the fitting process. The simultaneous fitting could also introduce unwanted interplay between source parameters and deadtime parameters, resulting in less accurate estimates for both. Jernigan et al. and Sunyaev & Revnivtsev do ensure that the best-fit values of the two deadtime coefficients A and B are comparable to their theoretical values. Nevertheless, their approach is not as satisfying when compared to an accurate, predictive model as an alternative.

In the case of Sunyaev & Revnivtsev (2000) in particular, this fitting method calls into question the validity of cut-off frequencies found after Poisson subtraction, since it supposes a priori that the power above a certain frequency is due entirely to Poisson statistics. The fitting process could create, somewhat artificially, a relative absence of high-frequency power in the region where it is applied. Sunyaev & Revnivtsev defined the fitting region to be $f \geq 50$ Hz for BHBs and $f \geq 600$ Hz for NSBs. The difference in the fitting regions for BHBs and NSBs could bias the comparison of source PDSs obtained after Poisson subtraction.

1.5 Outline of thesis

The remainder of the thesis is organized as follows: Chapter 2 briefly describes the physical characteristics and functioning of the RXTE PCA instrument. The different categories of events recorded by the RXTE PCA are summarized. Chapter 3 presents

the mathematical framework that governs the power spectral analysis of X-ray timing data, including a discrete-time treatment that is more representative of signal processing practices in observational astronomy. In Chapter 4, the new deadtime model is derived, building upon the model of Zhang et al. (1995) and incorporating the deadtime caused by non-selected events. The dependence of the model Poisson PDS on count rates and deadtime parameters is explored in a series of plots. Chapter 5 presents the results of testing the deadtime model using very-high-frequency PDSs obtained from RXTE observations. Measurements of the deadtime intervals t_d for selected events and τ_{VLE} for VLEs are described. Chapter 6 discusses known phenomena that were not incorporated into the deadtime model of this thesis. Deadtime corrections for highly non-stationary sources and deadtime effects dependent on the energies of events are of particular importance for further modelling efforts. Finally, in Chapter 7, we discuss a preliminary comparison between the PDSs of BHBs and NSBs using the new deadtime model. The work of Sunyaev & Revnivtsev (2000) is extended by considering the dependence of PDS characteristics, in particular the high-frequency continuum power, on the BHB X-ray state.

Chapter 2

The Rossi X-ray Timing Explorer Proportional Counter Array

The PCA aboard RXTE is a satellite-based X-ray observatory that offers fast, micro-second-level time resolution and modest spectral resolution at 2–60 keV for celestial X-ray sources. The RXTE PCA has been described in detail in a number of papers (e.g. Jahoda et al. 2006; Swank 1998). This chapter summarizes the essential functions and characteristics of the PCA that are relevant to constructing a model of PCA deadtime effects on the PDS.

2.1 Physical description

The PCA consists of 5 proportional counter units (PCU) collecting data in parallel. Briefly, a proportional gas counter is a device used to detect X-rays and measure their energies. Its main component is a gas chamber with a window transparent to X-rays on the top surface. The gas mixture is typically dominated by an inert noble gas with high atomic number, such as xenon. An X-ray that enters the chamber through the window is absorbed by a gas molecule and liberates a lower-shell electron. The electron is then accelerated toward a high-voltage anode to be collected. While being accelerated, the initial electron collides with other gas molecules and liberates additional electrons, resulting in an "avalanche." The number of subsequent electrons

released depends on the energy of the initial electron, which is in turn directly dependent on the energy of the incident photon. The overall effect of the X-ray induced electron avalanche is to deposit a charge on the anode proportional to the incident X-ray energy. By measuring the charge deposited by each photon and calibrating the proportional counter with line-emitting X-ray sources in the laboratory, incident photons can be counted and their energies determined. The RXTE PCA can detect X-ray events in the energy range 2–60 keV, where the lower limit is set by the transmission of the mylar window and the upper limit is set by the path length of X-rays in xenon in relation to the size of the detector.

Inevitably, the proportional counter will also admit background particles, which cause ionization events that need to be screened from photon events. Particles tend to exhibit a different ionization profile, leaving behind a track of multiple ionizations. The multiple ionizations allow the X-ray detector to distinguish and exclude particles from photons. The filtering is achieved through an anti-coincidence mechanism, which disqualifies events from different measurement chains that arrive too close to one another in time. The anti-coincidence mechanism will be described in greater detail later on in Section 2.2.

A schematic diagram for a single PCU is shown in Figure 2-1. Each PCU is topped with a hexagonally-shaped mechanical collimator about 20 cm long that exposes the PCU to a narrow sector of the sky with FWHM $\approx 1°$. The effective collecting area of a PCU is 1200 cm^2. The PCA does not form images; a source is observed by pointing the array at it and collecting all events that pass through the collimators. Thus, the PCA is most suitable for sources that are well-localized and bright compared to the cosmic X-ray background. In contrast, some X-ray observatories employ focusing telescopes that permit studies of very faint X-ray sources. However, detectors with imaging telescopes are limited to low event throughput. X-ray timing analyses would thereby suffer from poorer statistics as compared to a non-imaging instrument like the RXTE PCA.

The top window of a PCU is constructed of mylar coated on both sides with aluminum. The PCUs are housed in aluminum bodies shielded with tin and tantalum
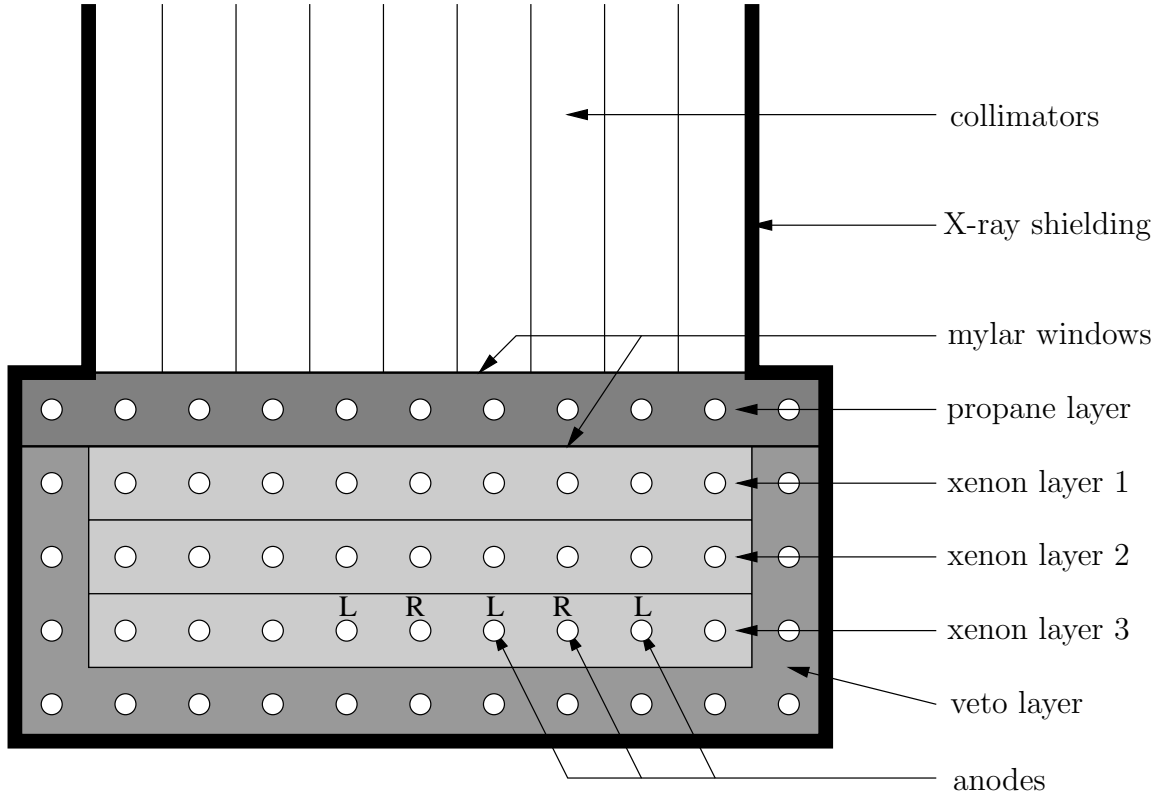
Figure 2-1: Schematic diagram of a single PCU.

so that all events must enter through the top window. Below the top window, the first collecting layer is a propane-filled filtering layer about 1 cm thick and equipped with its own collecting anode. Particles and other non-photon events interact with the propane and lose some to all of their energy, while most X-rays pass through unabsorbed. In this manner, the xenon-filled main counter, below the propane layer and separated by another aluminum-coated mylar window, is partially protected from non-photon events.

The main xenon chamber is divided into collecting cells 1 m long, 1 cm wide and 1 cm thick. The cells are bounded by wire anodes and are arranged into three vertical layers numbered 1, 2, and 3. The cell wires in a layer are connected either to the layer's left or right chain in an alternating fashion, resulting in six chains 1L, 1R, 2L, 2R, 3L, 3R that fill the xenon chamber.

Surrounding the main xenon chamber is an additional xenon veto layer that is shielded from a direct view of the X-ray sky. Events that trigger the xenon veto layer

are likely particle events and hence are excluded.

In sum, each PCU produces 7 outputs, 6 from the xenon half-layers and 1 from the propane layer. Each of the 7 outputs is fed into its own analog electronics chain consisting of a charge-sensitive amplifier, a shaping amplifier, a discriminator and a peak detector. Together these analog circuits integrate the charge associated with an event, convert the charge into an equivalent voltage pulse, and determine the peak of the pulse. The output voltage from each analog chain thus indicates the energy of the event it received.

As can be seen in Figure 2-2, the 7 amplifier chains are reunited so as to share a single analog-to-digital converter (ADC). The ADC is the slowest part of the measurement chain by a substantial margin. After a nominal conversion time of $t_d = 10$ μs, the ADC outputs an 8-bit digital word to represent the energy of the incident event. The essential consequence of sharing a single ADC among 7 analog chains is that the ADC becomes unavailable to all chains once it is engaged by an event on any one chain. For example, if an event travels through the 1L chain and triggers the ADC, events on *all* chains that arrive at the ADC within the next t_d microseconds are lost. The 10 μs conversion time is a non-paralyzable deadtime because an ADC does not pay any attention to its input while it is busy converting a latched voltage, so there is no means by which subsequent events can extend the deadtime.

2.2   Types of events

We distinguish three types of events that initiate analog-to-digital conversions. The first type consists of the "good" events that trigger only a single xenon chain and are most likely caused by X-rays. The second type consists of "propane" events, events that trigger only the propane chain. The third type will be referred to as "bad" events and includes coincident events that trigger multiple anodes. Coincident events are likely particle events and thus are not included among the good events. As defined by the detector, two or more events are considered coincident when they all occur within an anti-coincidence window of length $3-4$ μs *and* trigger more than one
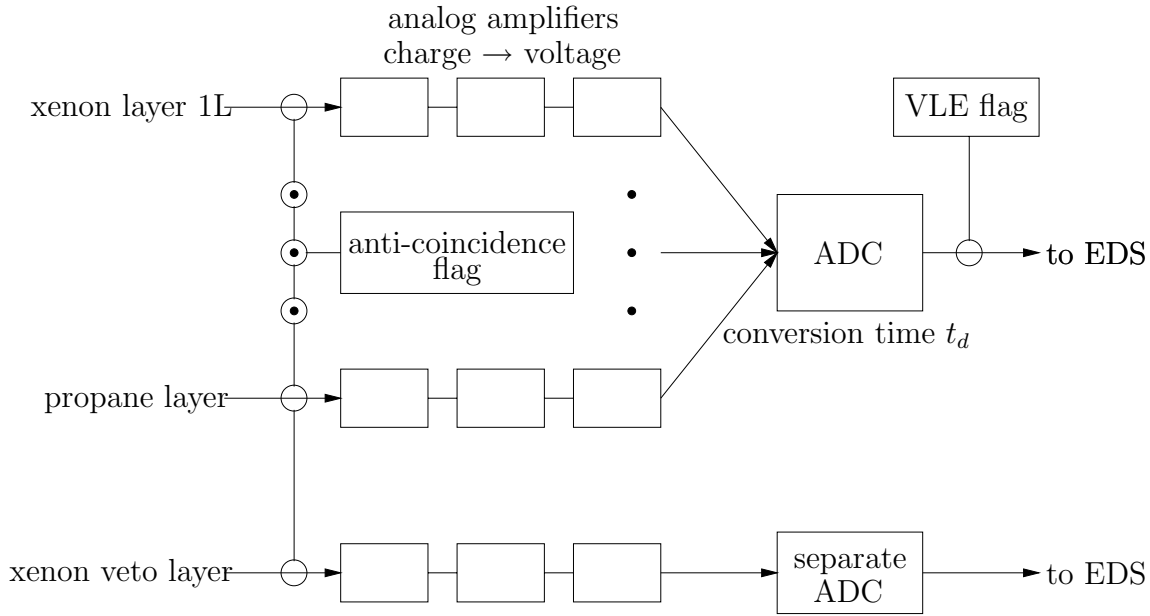
Figure 2-2: Measurement chains for a single PCU.

measurement chain. Two or more events that occur within a single anti-coincidence window but trigger only one chain cannot be distinguished from a single event, as the pulses tend to overlap and yield an inflated energy measurement. The anti-coincidence mechanism is represented in Figure 2-2 by the "anti-coincidence flag" connected to all of the measurement chains.

Bad events can also arise when the xenon veto layer is triggered, either alone or in conjunction with other anodes. However, when the xenon veto layer is triggered alone, a separate 2-bit ADC is employed and the main ADC is not engaged. The bad event count recorded in Standard 1 timing mode data (to be discussed at the end of this section) over-estimates the number of bad events that actually initiate analog-to-digital conversions by the number of xenon veto layer-only events. The rate of xenon veto layer-only events is less than 10 counts/sec/PCU (Morgan 2006).

Only the good events are relevant for scientific observations. However, the sharing of the ADC means that the propane and bad events can interfere with observations of good events. Propane events can engage an idle ADC and prevent it from registering good events. Likewise, the first event of a coincident set will appear to be a good event (or a propane event) and will trigger the ADC before being labelled "bad" upon

33

the arrival of the second event of the coincident set. The ADC is nonetheless busy for a time t_d following the first event of the coincident set.

VLEs constitute a fourth type of event. As mentioned in Section 1.4, VLEs are energetic particle events (> 120 keV equivalent X-ray energy for the current PCA gain setting) that can saturate the analog amplifiers and initiate a long, oscillatory relaxation toward the baseline. The oscillations brought on by saturation are often mistaken for actual events. To prevent the recording of such "ghost" events and interference with subsequent actual events, the detector is disabled following a VLE for an amount of time called the VLE window and denoted τ_{VLE}. No events that enter during a VLE window are recorded, although the VLEs themselves are counted. Two VLE windows of nominal lengths 70 and 170 μs are available to be selected by the observer (Jahoda et al. 2006). We note that values of 61/142 μs (Jernigan et al. 2000; Revnivtsev et al. 2000) and 55/155 μs have also been reported.

Every event in these four categories generates a transmission to the Experiment Data System (EDS) that encodes the event energy and the measurement chains that were affected. The EDS tags events with their arrival times and performs event selection and data compression according to user specifications. The EDS for the PCA consists of six event analyzers (EAs). Four EAs select data according to telemetry compression modes specified by the observer in order to best accommodate the particular scientific goals. The remaining two EAs routinely operate in Standard 1 timing and Standard 2 spectral modes for all observers. Of particular interest is the Standard 1 timing mode, which records with 1/8-second time resolution the number of good events for each PCU, and the numbers of propane, bad, and VLEs for the entire PCA. Standard 1 data is therefore very useful in modelling and quantifying the self-deadtime effects among selected events, as well as the interference effects caused by VLEs, propane, and bad events.

Chapter 3

X-ray Timing Analysis and Power Density Spectra

In this chapter we present the general framework for timing analysis in X-ray astronomy. Within this framework, we will then model the effects of detector deadtime on the PDS. Throughout this chapter, we make use of concepts and notation from Oppenheim et al. (1999), a widely-used signal processing textbook, as well as from papers that analyze deadtime effects in an astrophysical context (Zhang et al. 1995; Vikhlinin et al. 1994).

3.1 Light curves, random processes, and power density spectra

In continuous time, X-ray sources are described by a light curve $x(t)$, a function of time with dimensions [counts/second]. We may represent photon arrivals, which are discrete events, by Dirac delta functions of unit amplitude at their times of occurrence (Vikhlinin et al. 1994). The function $x(t)$ is then composed of a string of unit delta functions marking the arrival of events, as represented by Figure 3-1. A greater density of delta functions in an interval means that more events are arriving during that time.
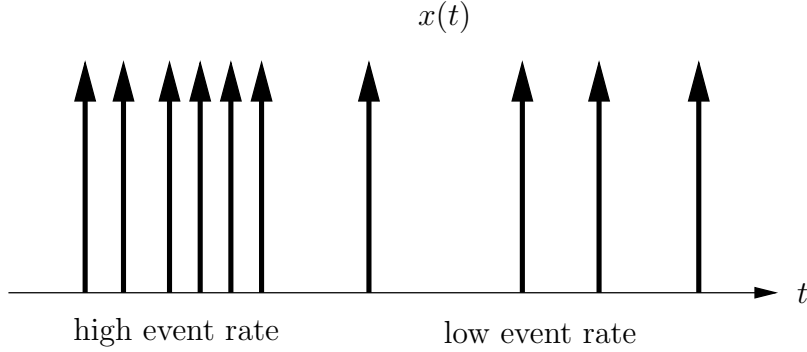
Figure 3-1: Light curve example $x(t)$.

We assume that the arrival process described by $x(t)$ is a wide-sense stationary, ergodic random process. The ergodic property allows us to define the mean and the autocorrelation function as time averages instead of ensemble averages (Oppenheim et al. 1999):

$$\mu_x = \lim_{T \to \infty} \frac{1}{2T} \int_{-T}^{T} x(t)\, dt \equiv r_o, \tag{3.1}$$

$$\phi_{xx}(\tau) = \lim_{T \to \infty} \frac{1}{2T} \int_{-T}^{T} x(t)x(t+\tau)\, dt. \tag{3.2}$$

Because of stationarity, the mean observed event rate r_o is independent of time, while the autocorrelation function $\phi_{xx}(\tau)$ is only a function of time lag τ.

The PDS is the Fourier transform of the autocorrelation function. Since autocorrelations are even functions of τ, $\phi_{xx}(\tau) = \phi_{xx}(-\tau)$, the PDS $P_{xx}(f)$ can be written as

$$P_{xx}(f) = 2 \int_{0}^{\infty} \phi_{xx}(\tau) \cos(2\pi f \tau)\, d\tau, \tag{3.3}$$

where f is the frequency in Hz. Equivalently, the autocorrelation function $\phi_{xx}(\tau)$ is proportional to $x(t)$ convolved with its time reversal $x(-t)$. For real $x(t)$, the PDS $P_{xx}(f)$ is then proportional to the magnitude squared of the Fourier transform $X(f)$ of $x(t)$:

$$P_{xx}(f) = \lim_{T \to \infty} \frac{|X(f)|^2}{2T}, \tag{3.4}$$

36

where the Fourier transform $X(f)$ (if it exists) is defined as

$$X(f) = \lim_{T \to \infty} \int_{-T}^{T} x(t) e^{-i2\pi ft} \, dt. \tag{3.5}$$

3.2 Light curves and power density spectra in observational astronomy

Astronomical measurements of X-ray sources occur in discrete time, not continuous time. We do not measure $x(t)$, but instead we count the number of events occurring in successive time bins of size t_b. The number of counts $x[n]$ in the nth time bin is related to $x(t)$ by

$$x[n] = \int_{nt_b}^{(n+1)t_b} x(t) \, dt. \tag{3.6}$$

For the purposes of analysis, the sequence of counts $x[n]$ is usually broken into non-overlapping segments $x_r[n]$ of length N:

$$x_r[n] = \begin{cases} x[n + rN], & 0 \leq n \leq N - 1, \\ 0, & \text{otherwise.} \end{cases} \tag{3.7}$$

For each segment, the mean number of counts per bin is estimated from the sample mean:

$$\langle x_r[n] \rangle = \frac{1}{N} \sum_{n=0}^{N-1} x_r[n] = r_o t_b. \tag{3.8}$$

The sample mean is then subtracted from $x_r[n]$ to yield a zero-mean sequence. Otherwise, a non-zero mean would create a large zero-frequency component in the spectrum, an uninteresting feature that could obscure nearby low-frequency components because of leakage due to windowing (Oppenheim et al. 1999). The spectral content of each segment $x_r[n]$ (now zero-mean) is found by taking its discrete Fourier transform (DFT):

$$X_r[k] = \sum_{n=0}^{N-1} x_r[n] e^{-i(2\pi/N)kn}. \tag{3.9}$$

To take advantage of the Fast Fourier Transform (FFT) algorithms, the segment length N is usually a power of 2.

From each segment we form an estimate of the continuous-time PDS $P_{xx}(f)$:

$$I_r[k] = \frac{|X_r[k]|^2}{Nt_b}. \tag{3.10}$$

The estimate $I_r[k]$ is known as the periodogram (Oppenheim et al. 1999). The time bin size t_b is included in the denominator to make $I_r[k]$ dimensionally equivalent to $P_{xx}(f)$.

While the periodogram estimate $I_r[k]$ from each segment does approach the true PDS $P_{xx}(f)$ in an average sense, its statistical uncertainty is generally very large. Therefore, in practice it is necessary to average periodograms from a large number of segments R, resulting in the average periodogram estimate of the PDS:

$$\langle I[k] \rangle = \frac{1}{R} \sum_{r=0}^{R-1} I_r[k]. \tag{3.11}$$

In X-ray astronomy, periodogram averaging is the usual method of estimating the PDS from a sequence of event counts. Strictly speaking, the average periodogram is only an estimate of the true PDS. However, while using this technique, observers assume that the two are equal.

3.3 Power density spectrum normalizations

Two different normalizations of the periodogram in equation (3.10) are popular in X-ray astronomy. The first is known as Leahy normalization (Leahy et al. 1983) and includes an additional factor of $2/r_o$:

$$I_r[k] = \frac{2|X_r[k]|^2}{Nr_ot_b} = \frac{2|X_r[k]|^2}{N_{ph}}, \tag{3.12}$$

where we recognize that the quantity Nr_ot_b is equal to the total number of events N_{ph} observed during the segment $x_r[n]$. The Leahy normalization has the property

that the PDS computed for a pure Poisson arrival process has the constant value 2 over the entire frequency band. Prior to Leahy normalization, the PDS for a pure Poisson process is still constant or white, since there is no event-to-event correlation at any frequency. However, the constant level would be proportional to the variance in the number of counts per time bin, which in turn is proportional to the observed rate r_o. Multiplying by $2/r_o$ removes the dependence on r_o.

The second commonly used normalization is called root-mean-square (RMS) normalization, in which the PDS is expressed in units of (RMS variability / mean)^2 per Hz. First, the mean observed rate r_o is corrected for the contribution from background (i.e. non-source) events, so that the corrected rate r'_o is a truer estimate of the source brightness. The RMS-normalized PDS is the right-hand side of equation (3.10) multiplied by $2/r'^2_o$:

$$I_r[k] = \frac{2|X_r[k]|^2}{N r'^2_o t_b}. \tag{3.13}$$

In X-ray astronomy, one often estimates the Leahy-normalized PDS using equations (3.12) and (3.11), subtracts the deadtime-corrected Poisson power level, and then applies RMS normalization. The resulting RMS-normalized source PDS is interpreted as the source variability per unit frequency as a fraction of the mean source rate r'_o. The integral of the RMS-normalized source PDS between frequencies f_1 and f_2 is equal to the contribution to the total variance from the frequency band $[f_1, f_2]$, normalized by the squared source rate r'^2_o. The RMS-normalized source PDS thus measures the magnitude of intrinsic source fluctuations relative to the mean source rate, an appealing interpretation for scientific analyses. Using RMS normalization, the variability of sources can be meaningfully compared even if the sources differ in brightness and the detectors used to observe them differ in sensitivity.

3.4 Relationship between the average periodogram and the discrete-time autocorrelation function

In using periodogram averaging to estimate the PDS, the discrete-time autocorrelation $\phi_{xx}[m]$ of $x[n]$ does not need to be calculated. However, for the purposes of modelling it is useful to express the average periodogram in equation (3.11) in terms of $\phi_{xx}[m]$, which we will define shortly. In doing so, we follow the approach of Zhang et al. (1995) closely. The reason for seeking such a relationship is as follows: We suppose that we can model the light curve $x(t)$ and derive its autocorrelation function. The PDS corresponding to the model can then be obtained from the autocorrelation function using the relationship between them.

We use the definition of the DFT in equation (3.9) and expand the Leahy-normalized PDS in equation (3.12),

$$I_r[k] = \frac{2}{N_{ph}} \sum_{n=0}^{N-1} \sum_{m=0}^{N-1} x[n]x[m]e^{-i(2\pi/N)k(n-m)}. \tag{3.14}$$

Next, we let R become very large so that the Leahy-normalized average periodogram $\langle I[k] \rangle$ given by equation (3.11) becomes an average over all time. Using ergodicity, $\langle I[k] \rangle$ can be interpreted as an ensemble average, so that $\langle I[k] \rangle$ is also equal to the ensemble average of equation (3.14). The result is an expression for the average periodogram in terms of the discrete-time autocorrelation $\phi_{xx}[m]$:

$$\langle I[k] \rangle = \frac{2}{r_o t_b}\left[\phi_{xx}[0] + 2\sum_{m=1}^{N-1} \frac{N-m}{N}\phi_{xx}[m]\cos\left(\frac{2\pi km}{N}\right) \right]. \tag{3.15}$$

The discrete-time autocorrelation function is defined as $\phi_{xx}[m] = E\left\{ x[n]x[n+m] \right\}$ and is also stationary (not a function of n). We use the notation $E\{\cdot\}$ to denote the expectation over an ensemble.

The relation between the discrete-time autocorrelation $\phi_{xx}[m]$ and its continuous-

time counterpart $\phi_{xx}(\tau)$ is as follows:

$$
\begin{aligned}
\phi_{xx}[m] &= \int_0^{t_b} d\tau_1 \int_{mt_b}^{(m+1)t_b} d\tau_2\ \phi_{xx}(\tau_2 - \tau_1), \\
&= t_b \int_{(m-1)t_b}^{(m+1)t_b} d\tau\ \phi_{xx}(\tau) \left(1 - \left| \frac{\tau}{t_b} - m \right| \right).
\end{aligned}
\tag{3.16}
$$

The second, simpler equality can be derived by a change of variables in the double integration. The relation can be thought of as multiplying $\phi_{xx}(\tau)$ by a triangular window of width $2t_b$, height t_b, and centre mt_b, and integrating the product.

In summary, given a model for the signal $x(t)$, we can derive its continuous-time autocorrelation $\phi_{xx}(\tau)$, convert $\phi_{xx}(\tau)$ to $\phi_{xx}[m]$ using equation (3.16), and compute the predicted PDS using equation (3.15). In Chapter 4 we will use this approach to derive the expected PDS for an initially Poisson arrival process that undergoes detector deadtime effects.

Chapter 4

PCA Deadtime Model

4.1   Overview

As discussed in Section 1.1, the PDS is employed to examine the variability of X-ray sources from measurements of the X-ray flux. However, we cannot observe the X-ray flux continuously, but rather must rely on discrete counts of X-ray events per time bin. Such counting measurements do yield the X-ray flux on average, but inevitably include Poisson statistical variations about the true event rate.

Taking the simplest example, if the X-ray source has a constant rate of r counts per second for all time, its constituent photons will follow a Poisson arrival process characterized by the rate r. The number of counts in any one-second interval will on average deviate from the mean r by the standard deviation $\sqrt{r}$. Because of this statistical variation, the resulting PDS will be non-zero. The PDS will also be white; since different times in a Poisson process are independent, the count data will show no correlation on any timescale.

Real X-ray sources however are never truly stationary, but exhibit an event rate that fluctuates about a local mean. The count data will contain variations due to the evolution of the source rate in addition to Poisson statistical variations. In the power spectral domain, the frequency content of the source rate variation will appear

in addition to the flat Poisson PDS:

$$P_{total}(f) = P_{source}(f) + P_{Poisson}(f). \tag{4.1}$$

In the earlier example of a source rate that is constant, the source power $P_{source}(f)$ is zero because the source flux does not vary about the mean rate r (and assuming mean subtraction), and the total power $P_{total}(f)$ is completely accounted for by the Poisson component $P_{Poisson}(f)$.

Our scientific interest lies not in Poisson statistics but in the variability $P_{source}(f)$ of the source flux. Therefore, it is common practice in analyzing PDSs to subtract the Poisson power level $P_{Poisson}(f)$, which is calculated by assuming a Poisson process with rate equal to the mean observed rate. The residual PDS $P_{total}(f) - P_{Poisson}(f)$ should then be a direct estimate of the intrinsic flux variability.

In the presence of detector deadtime, and specifically self-deadtime among selected events, the PDS $P_{Poisson}(f)$ of a stationary Poisson process is no longer flat versus frequency. Qualitatively speaking, detector deadtime eliminates any correlations in the count data at time lags less than the deadtime, since it is impossible for two events to be separated by less than the deadtime. This gap in the autocorrelation function affects the PDS. The modified PDS will depend upon the deadtime characteristics of the detector in addition to the average arrival rate. It is the modified Poisson PDS that must now be subtracted from the total PDS. Thus, it is necessary to derive the precise form of the modified Poisson PDS to ensure an accurate subtraction and estimation of the intrinsic variability.

For the RXTE PCA instrument, and for proportional counters in general, the deadtime is dominated by the ADC conversion time. In addition to the self-deadtime among selected events, interference from non-selected events that compete for ADC measurement will also impact the autocorrelation and the PDS. Non-selected events introduce additional gaps where no good events can be observed.

The remainder of this chapter will be concerned with deriving the PDS $P_{Poisson}(f)$ corresponding to a Poisson arrival process observed through a detector that features

both self-deadtime and interference from non-selected events. Figure 4-1 shows the model of the deadtime and discretization processes that take place inside the detector. The stages of the derivation will follow the processing stages in Figure 4-1.
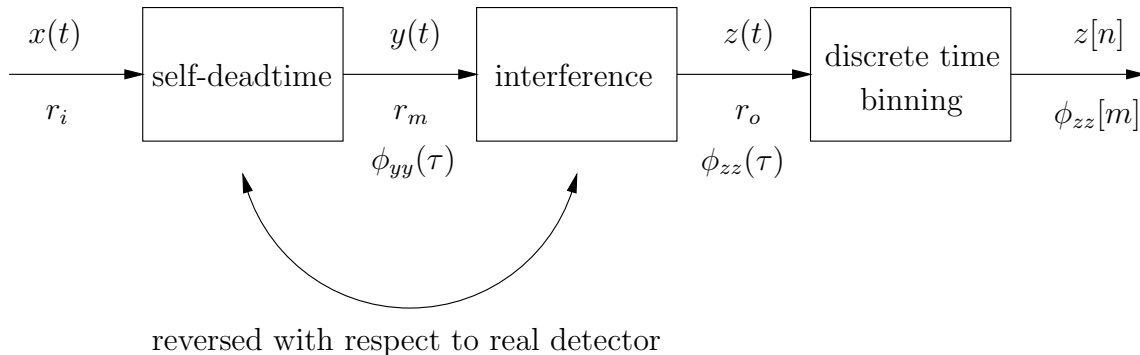

Figure 4-1: Block diagram for deadtime processes in the RXTE PCA.

We should point out a key simplification made in the present model. In the actual PCA, good events first undergo an interference process where they may be blocked from the ADC by VLEs, propane, and "bad" events. The self-deadtime effect among good events can happen only after a good event avoids being blocked. In other words, one can think of the interference process as occurring first, followed by the self-deadtime process. However, the temporal correlations introduced by the interference process cause the input into the self-deadtime process to lose the simplifying memorylessness property of Poisson processes. An analysis of the self-deadtime process with a non-Poisson input becomes significantly more complicated (see Section 6.4). Therefore, in the interests of obtaining a simpler mathematical model, we reverse the order of the two processes. As long as the losses due to deadtime and interference are small compared to the overall flux, the inaccuracy arising from the reversal should be unimportant.

Let the light curve $x(t)$ represent "good" X-ray events that arrive at the detector according to a Poisson process. $x(t)$ includes all single (i.e. non-coincident) events that pass through one of the six xenon chains and is characterized by a constant rate of r_i events per second per PCU.[1] We suppose that $x(t)$ first undergoes a self-

45

deadtime process, losing some events to deadtime to yield the output $y(t)$. $y(t)$ is then subjected to interference by non-selected events, resulting in the observed light curve $z(t)$.

We seek the observed event rate r_o and the continuous-time autocorrelation function $\phi_{zz}(\tau)$ for $z(t)$. r_o and $\phi_{zz}(\tau)$ may be derived from r_m and $\phi_{yy}(\tau)$, the event rate and autocorrelation function associated with the output $y(t)$ of the self-deadtime process. Once $\phi_{zz}(\tau)$ is obtained, equation (3.16) converts $\phi_{zz}(\tau)$ to its counterpart $\phi_{zz}[m]$ in discrete binned time. The PDS is then determined from $\phi_{zz}[m]$ using equation (3.15).

4.2   Self-deadtime among good events

As previously mentioned, the self-deadtime effect is primarily caused by the analog-to-digital conversion of good events and is non-paralyzable in nature. We denote the conversion time by t_d. Zhang et al. (1995) have derived the output event rate r_m and the continuous-time autocorrelation $\phi_{yy}(\tau)$ for a non-paralyzable deadtime process:

$$r_m = \frac{r_i}{1 + r_i t_d}, \tag{4.2}$$

$$\phi_{yy}(\tau) = r_m \delta(\tau) + r_i r_m \sum_{k=1}^{\infty} u(x_k) \frac{(r_i x_k)^{k-1}}{(k-1)!} e^{-r_i x_k}, \tag{4.3}$$

where $x_k = \tau - k t_d$ and $u(x_k)$ is the unit step function: $u(x_k) = 1$, $x_k \geq 0$ and $u(x_k) = 0$, $x_k < 0$. Equation (4.2) states that the average interval $1/r_m$ between detected events is the sum of the deadtime t_d and $1/r_i$, the expected time from the end of the conversion until the next arrival. To derive equation (4.3), Zhang et al. consider the probability of detecting an event at lag τ relative to the previous event given a non-paralyzable deadtime of t_d. We note that $\phi_{yy}(\tau) = 0$ for $0 < \tau < t_d$, as is expected since events cannot be separated by less than t_d.

---

per PCU is unavailable, the rate for the entire PCA is divided by the number of active PCUs.

4.3 Interference from non-selected events

Armed with the mean event rate r_m and the autocorrelation function $\phi_{yy}(\tau)$ for the X-ray flux $y(t)$ emerging from the self-deadtime process, we now consider the interference of VLEs, propane events and bad events. If X-rays from only a single energy band have been selected, good events in other energy bands are also considered to be non-selected events because they prevent the detection of X-rays in the selected band.

We can represent the result of the interference process as $z(t) = v(t)y(t)$, where $v(t)$ is the masking function associated with the non-selected events. $v(t)$ is equal to zero whenever the detector is occupied with a non-selected event, and is equal to unity otherwise. More precisely:

- $v(t) = 0$ for intervals of length τ_{VLE} that occur at a rate of r_{VLE} intervals per second. These intervals represent the disabling of the detector for a non-paralyzable time τ_{VLE} following VLEs, which occur at a rate r_{VLE}.

- $v(t) = 0$ for intervals of length t_d that occur at a rate r_1, equal to the sum of the propane rate r_p, the bad rate r_b, and the out-of-band good event rate r_{ob}.

- $v(t) = 1$ otherwise.

The masking function $v(t)$ is depicted in Figure 4-2.

The types of non-selected events identified above are considered to have disjoint probabilities. If we define r_{VLE}, r_p, r_b, and r_{ob} as the *observed* rates recorded in the Standard 1 data mode, then the corresponding non-selected events cannot overlap each other. Each recorded non-selected event, because it was recorded, must have initiated its own distinct masking interval. Thus, as long as we use observed rates of non-selected events, we need not be concerned with masking intervals overlapping each other.

We assume that the non-selected events represented in $v(t)$ are independent of the selected events in $y(t)$. In the case of VLEs and the large majority of propane and bad events, the events are caused by particles in the path of the satellite and are unrelated
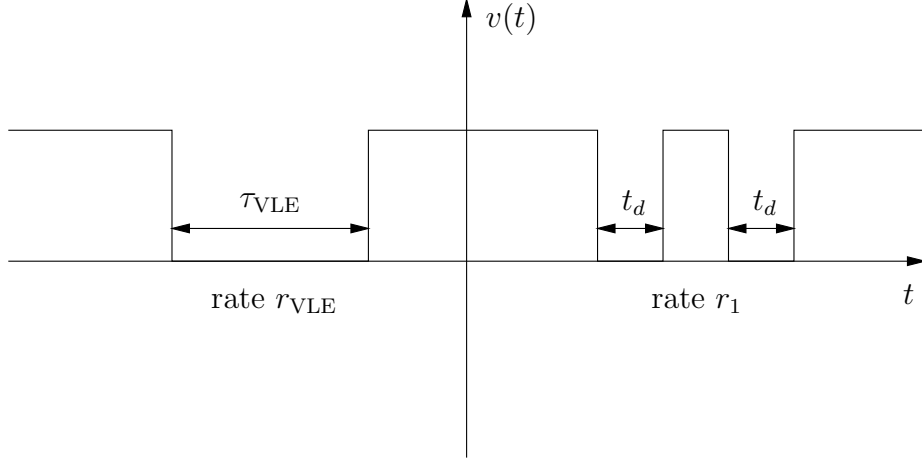
Figure 4-2: The masking function $v(t)$.

to X-ray photons from the source. On the other hand, multiple photon coincidences and out-of-band good events are correlated with the selected events, since an increase in overall flux increases the likelihood of coincidences and usually increases the flux in each energy band. Nevertheless, the assumption of independence between $v(t)$ and $y(t)$ is reasonable and greatly simplifying.

The observed rate r_o corresponding to $z(t)$ can thus be written as $r_o = r_m E\left\{v(t)\right\} = r_m \mu_v$, where μ_v denotes the expected value of $v(t)$. Similarly, the autocorrelation of $z(t)$ can be expressed as:

$$
\begin{aligned}
\phi_{zz}(\tau) = E\left\{z(t)z(t+\tau)\right\} &= E\left\{v(t)y(t)v(t+\tau)y(t+\tau)\right\}, \\
&= E\left\{v(t)v(t+\tau)\right\} E\left\{y(t)y(t+\tau)\right\}, \\
&= \phi_{vv}(\tau)\phi_{yy}(\tau). \tag{4.4}
\end{aligned}
$$

r_m and $\phi_{yy}(\tau)$ are given in equations (4.2) and (4.3); it remains to calculate μ_v and $\phi_{vv}(\tau)$, the autocorrelation function for the masking function $v(t)$. We have also assumed that $v(t)$ is wide-sense stationary, so that $\phi_{vv}(\tau)$ is a function of τ only and not both t and τ. The assumption of wide-sense stationarity is valid as long as the rates of non-selected events do not contain large secular variations over the course of an observation.

In calculating μ_v and $\phi_{vv}(\tau)$ it is somewhat easier to consider $w(t) = 1 - v(t)$

48

instead. Thus $w(t) = 1$ whenever a masking interval is in effect, $w(t) = 0$ otherwise, and $w(t)$ is also wide-sense stationary.

$$\mu_v = E\{1 - w(t)\} = 1 - \mu_w, \tag{4.5}$$

$$\phi_{vv}(\tau) = E\{(1 - w(t))(1 - w(t + \tau))\},$$

$$= 1 - E\{w(t)\} - E\{w(t + \tau)\} + E\{w(t)w(t + \tau)\},$$

$$= 1 - 2\mu_w + \phi_{ww}(\tau). \tag{4.6}$$

The expectation value μ_w is equal to the probability $Pr\{w(t) = 1\}$, since $w(t)$ is a binary process. $Pr\{w(t) = 1\}$ is in turn equal to the sum of the disjoint probabilities $r_i t_i$ for each type of non-selected event, where r_i and t_i are the associated rate and deadtime. Hence,

$$\mu_w = r_{\mathrm{VLE}}\tau_{\mathrm{VLE}} + (r_p + r_b + r_{ob})t_d = r_{\mathrm{VLE}}\tau_{\mathrm{VLE}} + r_1 t_d. \tag{4.7}$$

Similarly, the autocorrelation $\phi_{ww}(\tau)$ reduces to the following:

$$\phi_{ww}(\tau) = E\{w(t)w(t + \tau)\},$$

$$= Pr\{w(t)w(t + \tau) = 1\},$$

$$= Pr\{w(t) = 1\}Pr\{w(t + \tau) = 1|w(t) = 1\}, \tag{4.8}$$

where $Pr\{A|B\}$ means the probability of A given B. We divide the preceding joint probability into several cases based on the range of τ and the position of t relative to a masking interval:

Case $\tau \geq \tau_{\mathrm{VLE}}$: In this case, $w(t + \tau)$ is independent of $w(t)$ because any masking interval that was in effect at time t has expired before time $t + \tau$. Therefore $Pr\{w(t + \tau) = 1|w(t) = 1\} = Pr\{w(t + \tau) = 1\}$, which is the same as $Pr\{w(t) = 1\}$. From equation (4.7) we have:

$$\phi_{ww}(\tau) = (r_{\mathrm{VLE}}\tau_{\mathrm{VLE}} + r_1 t_d)^2, \quad \tau \geq \tau_{\mathrm{VLE}}. \tag{4.9}$$

Case $t_d \leq \tau < \tau_{\mathrm{VLE}}$: We break this into three sub-cases:

(a) time t falls into a masking interval of length t_d with probability $r_1 t_d$. Then $w(t+\tau)$ is independent of $w(t)$ because the interval containing t ceases to have any effect at $t + \tau$. At time $t + \tau$, all types of non-selected events can occur, not just those in the r_1 category, so $Pr\{w(t+\tau) = 1\} = r_{\mathrm{VLE}}\tau_{\mathrm{VLE}} + r_1 t_d$.

(b) time t falls within the first $\tau_{\mathrm{VLE}}-\tau$ of a VLE interval with length τ_{VLE}. Sub-case (b) occurs with probability $r_{\mathrm{VLE}}(\tau_{\mathrm{VLE}}-\tau)$. Then $Pr\{w(t+\tau) = 1|w(t) = 1\} = 1$ since $t + \tau$ is part of the same masking interval that contains t.

(c) time t falls within the last τ of a VLE interval with probability $r_{\mathrm{VLE}}\tau$. Then the interval that contains t ends before $t + \tau$ and $w(t+\tau)$ is again independent of $w(t)$ so that $Pr\{w(t + \tau) = 1|w(t) = 1\} = r_{\mathrm{VLE}}\tau_{\mathrm{VLE}} + r_1 t_d$.

The three sub-cases for case $t_d \leq \tau < \tau_{\mathrm{VLE}}$ are depicted in Figure 4-3.

Summing the probabilities from the three sub-cases,

$$\phi_{ww}(\tau) = r_1 t_d (r_{\mathrm{VLE}}\tau_{\mathrm{VLE}}+r_1 t_d)+r_{\mathrm{VLE}}(\tau_{\mathrm{VLE}}-\tau)+r_{\mathrm{VLE}}\tau(r_{\mathrm{VLE}}\tau_{\mathrm{VLE}}+r_1 t_d), \quad t_d \leq \tau < \tau_{\mathrm{VLE}}.$$
(4.10)

Case $0 \leq \tau < t_d$: We break this into four sub-cases:

(a) time t falls within the first $t_d - \tau$ of an interval of length t_d. Sub-case (a) occurs with probability $r_1(t_d - \tau)$. Then $Pr\{w(t + \tau) = 1|w(t) = 1\} = 1$ since $t + \tau$ is part of the same masking interval as t.

(b) time t falls within the last τ of an interval of length t_d with probability $r_1 \tau$. Then the interval that contains t ends before $t + \tau$ and $Pr\{w(t+\tau) = 1|w(t) = 1\} = r_{\mathrm{VLE}}\tau_{\mathrm{VLE}} + r_1 t_d$.

(c) time t falls within the first $\tau_{\mathrm{VLE}}-\tau$ of a VLE interval with probability $r_{\mathrm{VLE}}(\tau_{\mathrm{VLE}}-\tau)$. This is the same as case $t_d \leq \tau < \tau_{\mathrm{VLE}}$, sub-case (b).

(d) time t falls within the last τ of a VLE interval with probability $r_{\mathrm{VLE}}\tau$. This is the same as case $t_d \leq \tau < \tau_{\mathrm{VLE}}$, sub-case (c).
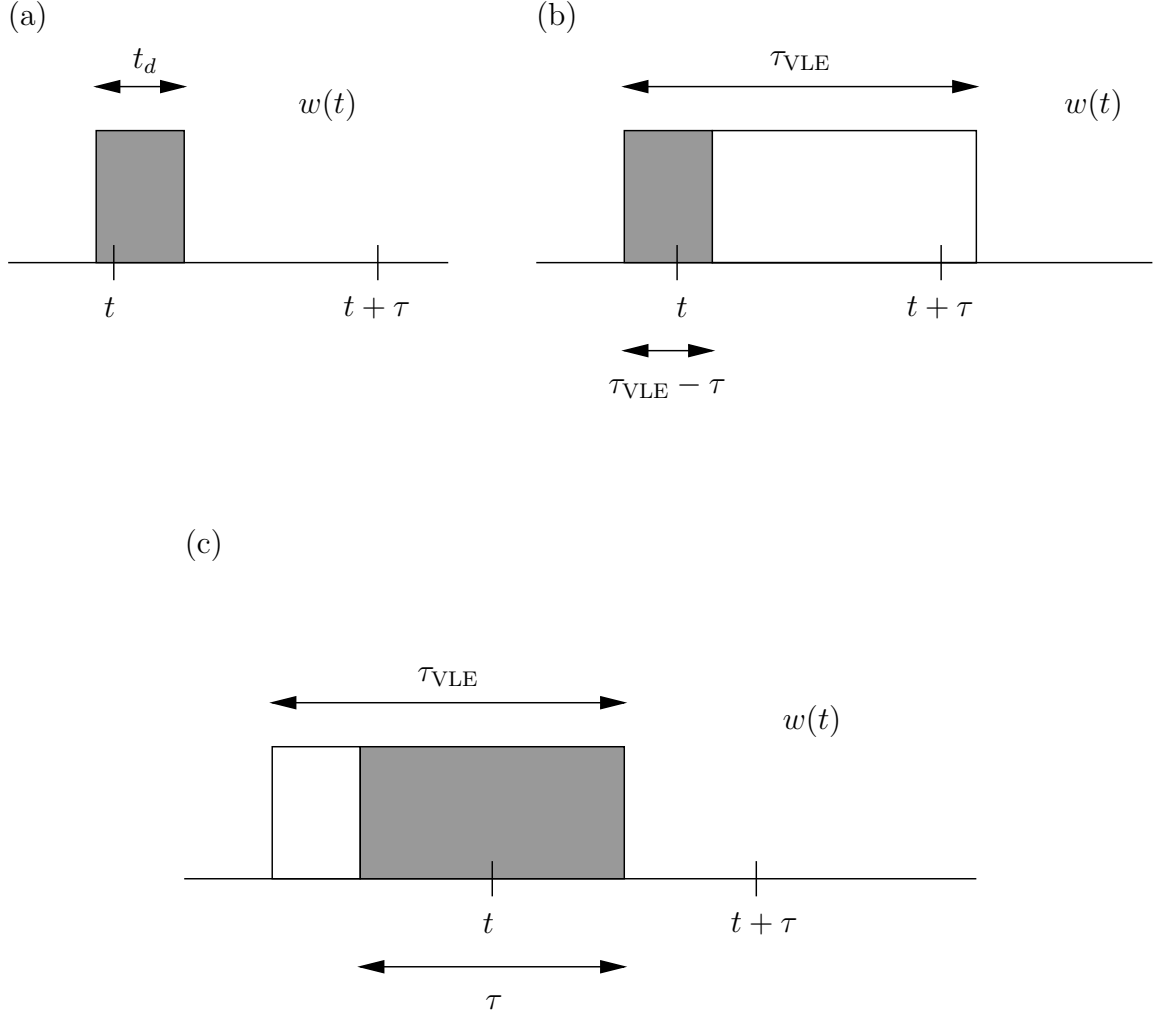
Figure 4-3: The three sub-cases for case $t_d \leq \tau < \tau_{\mathrm{VLE}}$.

The four sub-cases for case $0 \leq \tau < t_d$ are depicted in Figure 4-4.

Summing the probabilities from the four sub-cases,

$$\phi_{ww}(\tau) = r_1(t_d - \tau) + r_1\tau(r_{\mathrm{VLE}}\tau_{\mathrm{VLE}} + r_1 t_d) + r_{\mathrm{VLE}}(\tau_{\mathrm{VLE}} - \tau) + r_{\mathrm{VLE}}\tau(r_{\mathrm{VLE}}\tau_{\mathrm{VLE}} + r_1 t_d), \tag{4.11}$$

for $0 \leq \tau < t_d$.

We now compute μ_v and $\phi_{vv}(\tau)$ according to the relations in (4.5) and (4.6) and the quantities in equations (4.7) through (4.11).

$$\mu_v = 1 - r_1 t_d - r_{\mathrm{VLE}}\tau_{\mathrm{VLE}}. \tag{4.12}$$
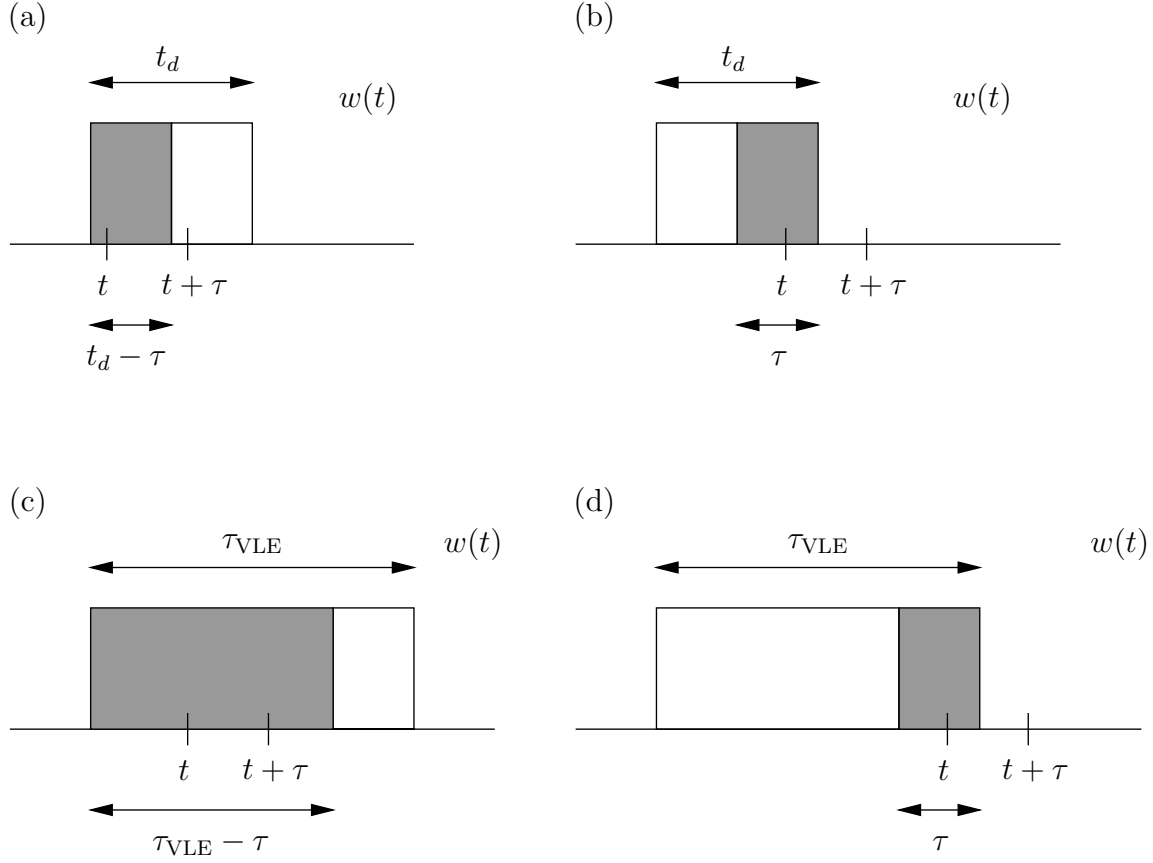
Figure 4-4: The four sub-cases for case $0 \leq \tau < t_d$.

Some factoring of terms yields the autocorrelation $\phi_{vv}(\tau)$:

$$\phi_{vv}(\tau) = \begin{cases} (1 - r_1 t_d - r_{\text{VLE}}\tau_{\text{VLE}})(1 - r_1\tau - r_{\text{VLE}}\tau), & 0 \leq \tau < t_d, \\ (1 - r_1 t_d - r_{\text{VLE}}\tau_{\text{VLE}})(1 - r_1 t_d - r_{\text{VLE}}\tau), & t_d \leq \tau < \tau_{\text{VLE}}, \\ (1 - r_1 t_d - r_{\text{VLE}}\tau_{\text{VLE}})^2, & \tau \geq \tau_{\text{VLE}}. \end{cases} \quad (4.13)$$

Finally, the observed event rate r_o and the autocorrelation function $\phi_{zz}(\tau)$ of $z(t)$ are given by

$$r_o = r_m(1 - r_1 t_d - r_{\text{VLE}}\tau_{\text{VLE}}) = \frac{r_i(1 - r_1 t_d - r_{\text{VLE}}\tau_{\text{VLE}})}{1 + r_i t_d}, \quad (4.14)$$

$$\phi_{zz}(\tau) = \begin{cases} r_o\,\delta(\tau), & 0 \leq \tau < t_d, \\[2ex] r_i r_o (1 - r_1 t_d - r_{\mathrm{VLE}}\tau) \displaystyle\sum_{k=1}^{\infty} u(x_k)\dfrac{(r_i x_k)^{k-1}}{(k-1)!}e^{-r_i x_k}, & t_d \leq \tau < \tau_{\mathrm{VLE}}, \\[3ex] r_i r_o (1 - r_1 t_d - r_{\mathrm{VLE}}\tau_{\mathrm{VLE}}) \displaystyle\sum_{k=1}^{\infty} u(x_k)\dfrac{(r_i x_k)^{k-1}}{(k-1)!}e^{-r_i x_k}, & \tau \geq \tau_{\mathrm{VLE}}, \end{cases}$$

$$(4.15)$$

where $x_k = \tau - k t_d$ as in equation (4.3). We have used equation (4.14) to substitute r_o for r_m. The incident event rate r_i can be inferred by solving equation (4.14) for r_i.

4.4 Computing the deadtime-corrected Poisson power density spectrum

Having derived the continuous-time autocorrelation function $\phi_{zz}(\tau)$ for a Poisson process with deadtime and interference effects, its discrete-time counterpart $\phi_{zz}[m]$ is determined by integrating $\phi_{zz}(\tau)$ using equation (3.16). As the expression for $\phi_{zz}(\tau)$ is rather complicated, the integration is accomplished numerically to the desired precision.

Calculating the PDS exactly from the discrete-time autocorrelation $\phi_{zz}[m]$ using equation (3.15) requires N values of $\phi_{zz}[m]$. However, we observe that $\phi_{zz}[m]$ approaches the constant value $r_o^2 t_b^2$ for large m (Zhang et al. 1995). This asymptotic value is in fact the square of the expected number of counts per time bin. Reaching the asymptotic value reflects the lack of correlation between counts in sufficiently separated time bins, so that mathematically we have $\phi_{zz}[m] = E\{z[n]z[n+m]\} \approx E\{z[n]\}E\{z[n+m]\} = r_o^2 t_b^2$.

Furthermore, we recall from Section 3.2 that the mean number of counts $r_o t_b$ is subtracted from measured data before forming the periodogram estimate. In order to be consistent with the periodogram method, the same mean subtraction is required here in computing the model PDS. In terms of the autocorrelation function, mean

53

subtraction is achieved by subtracting $r_o^2 t_b^2$ from $\phi_{zz}[m]$:

$$\tilde{\phi}_{zz}[m] = \phi_{zz}[m] - r_o^2 t_b^2. \tag{4.16}$$

As a consequence, $\tilde{\phi}_{zz}[m] \to 0$ as m increases.

The model PDS in equation (3.15) can now be written in terms of $\tilde{\phi}_{zz}[m]$. Let m_{max} be the value of m beyond which $\tilde{\phi}_{zz}[m]$ is negligible with respect to the desired accuracy, i.e. $|\tilde{\phi}_{zz}[m]| < \epsilon$, $m > m_{max}$ for some tolerance ϵ. Sacrificing a small amount of accuracy allows us to truncate the sum in (3.15) to m_{max}, which is usually considerably smaller than N:

$$\langle I[k] \rangle \simeq \frac{2}{r_o t_b} \left[\tilde{\phi}_{zz}[0] + 2 \sum_{m=1}^{m_{max}} \frac{N-m}{N} \tilde{\phi}_{zz}[m] \cos\left(\frac{2\pi km}{N}\right) \right] \tag{4.17}$$

We have now accomplished our goal of computing the PDS for a Poisson arrival process in the presence of deadtime and interference effects.

We employ the following procedure to subtract the deadtime-modified Poisson level using the new model from PDSs calculated from measured data: From the segment of measured data, we extract the observed rates r_o, r_1, and r_{VLE}. These rates, together with the timescales t_d and τ_{VLE}, parameterize the continuous-time autocorrelation $\phi_{zz}(\tau)$ in equation (4.15). $\phi_{zz}(\tau)$ is then evaluated on a dense set of points before being integrated numerically to obtain $\phi_{zz}[m]$. After subtracting the mean squared $r_o^2 t_b^2$ and truncating the sum as in equation (4.17), the model PDS is computed. The model PDS is then subtracted from the average periodogram estimate of the data, yielding an estimate of the source PDS.

4.5 Dependence of model power density spectrum on parameters

Before continuing on to applications of the deadtime model to real data, it is useful and illuminating to explore the dependence of the model on its parameters. These

parameters are: the observed rate of selected events r_o, the observed VLE rate r_{VLE}, the observed rate of other non-selected events r_1, the ADC deadtime t_d, the VLE window τ_{VLE}, the time bin size t_b, and the number of time bins per segment N. In this section, the dependence of the model PDS on the more important parameters is illustrated in a series of plots.

We first distinguish between the two regimes $t_b > t_d$ and $t_b < t_d$. In scientific observations, the time bin size t_b is usually chosen to yield a Nyquist frequency of a few kHz. Since $t_d \approx 10~\mu$s, the $t_b > t_d$ regime prevails for nearly all scientific observations. However, the RXTE PCA is capable of time resolutions down to 1 μs where $t_b < t_d$. Observations at such high time resolutions are undertaken to explore systematics of the detector, such as deadtime effects as will be described in Section 5.2.

In the case $t_b > t_d$, only the first few mean-subtracted autocorrelation values $\tilde{\phi}_{zz}[m]$ are significant. We expect that the correlations introduced by deadtime should die out as the lag τ becomes much longer than the deadtime timescales t_d and τ_{VLE}. Equivalently in discrete time, there should be little correlation at high values of m corresponding to lags $\tau = mt_b$. As is often the case, if $t_b >> t_d$ and $t_b \sim \tau_{\mathrm{VLE}}$, the correlations become negligible at relatively low values of m. Consequently, m_{max} in equation (4.17) is small and the PDS contains only a few low-frequency sinusoidal components. The PDS has the appearance of a half-wavelength of a sine wave superimposed on a constant level.

In contrast, if $t_b < t_d$, the correlation introduced by deadtime becomes sufficiently small only for large values of m. The upper bound m_{max} in equation (4.17) is large and the PDS is the sum of many sinusoidal components. The model PDS thus acquires a wavy appearance, resembling a sinc function subtracted from a constant level. The horizontal scaling is controlled by the ADC deadtime t_d. We might have expected this sinc-like behaviour from equation (1.1) (Zhang et al. 1995), putting aside the differences between paralyzable and non-paralyzable deadtime, and the continuous-time PDS versus the discrete-time PDS.

In keeping with the two regimes $t_b > t_d$ and $t_b < t_d$, the subsequent set of plots are

divided into two groups of four, the first group corresponding to $t_b > t_d$ and the last group to $t_b < t_d$. More precisely, $t_b = 2^{-13}$ s ≈ 122 μs for the first group and $t_b = 2^{-20}$ s ≈ 1 μs for the second group. In each group of four, the parameters being varied are, in order, r_o, t_d, r_{VLE}, and τ_{VLE}, while the other parameters are held constant. The number of time bins per segment is fixed at $N = 2^{20}$ throughout.

Figure 4-5 shows the dependence of the model PDS on the observed selected rate r_o, which ranges from 2500 counts/sec/PCU to 20000 counts/sec/PCU, for the regime $t_b > t_d$. The other parameters are fixed at $t_d = 10$ μs, $r_{\mathrm{VLE}} = 100$ counts/sec/PCU, $\tau_{\mathrm{VLE}} = 170$ μs (the nominal value for VLE setting 2), and $r_1 = 4000$ counts/sec/PCU. The values of r_{VLE} and r_1 represent typical values encountered in observing most celestial sources. The Poisson level in Leahy normalization decreases from around 1.9 to below 1.4 as r_o increases, while the curvature in the Poisson level also increases with r_o. We conclude that r_o controls the amplitude of the deadtime effect, as is expected since a higher selected count rate will incur more deadtime. Recall that the Leahy-normalized Poisson power should be a constant 2 in the absence of deadtime.

In Figure 4-6 we show the dependence on t_d, which was varied between 6 μs and 14 μs, for $t_b > t_d$. The other parameters are held at $r_o = 10000$ counts/sec/PCU, $r_{\mathrm{VLE}} = 100$ counts/sec/PCU, $\tau_{\mathrm{VLE}} = 170$ μs, and $r_1 = 4000$ counts/sec/PCU. Again we observe that the Poisson level decreases as t_d increases. However, in this case the amount of curvature varies inversely with respect to t_d, decreasing from top to bottom.

Figure 4-7 shows the model PDSs obtained for three different values of r_{VLE} and $t_b > t_d$. The other parameters are held constant at $r_o = 10000$ counts/sec/PCU, $t_d = 10$ μs, $\tau_{\mathrm{VLE}} = 170$ μs, and $r_1 = 4000$ counts/sec/PCU. For the range of VLE rates encountered by the RXTE PCA in practice, the effect due to r_{VLE} is more subtle than those due to r_o and t_d. r_o and t_d can be considered the primary parameters in the deadtime model while r_{VLE} and τ_{VLE} play secondary roles. Nevertheless, by increasing r_{VLE}, the lowest frequencies in the PDS are emphasized by a few hundredths to a tenth in Leahy units.
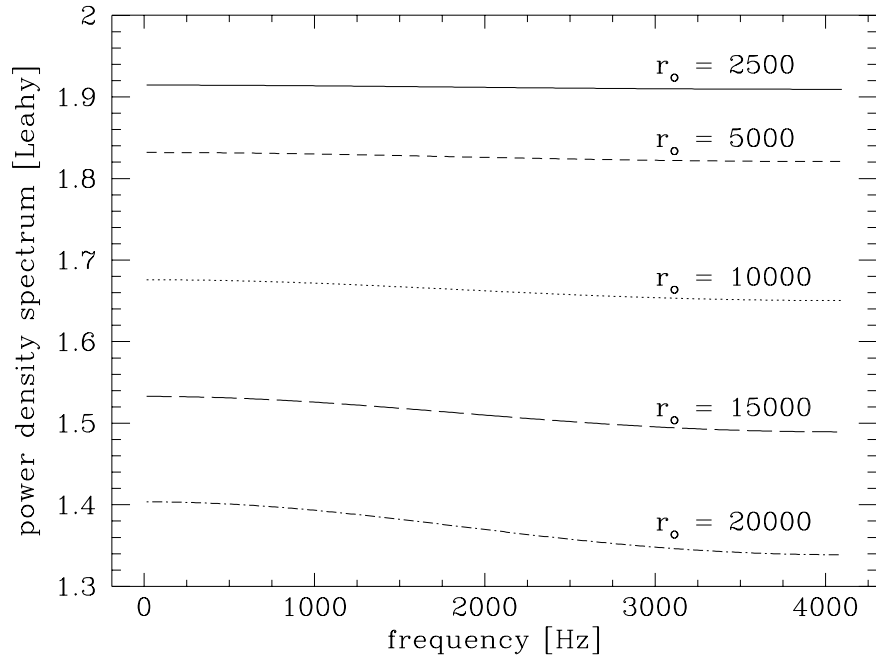
Figure 4-5: Model PDS for different values of r_o [counts/sec/PCU] and $t_b > t_d$.
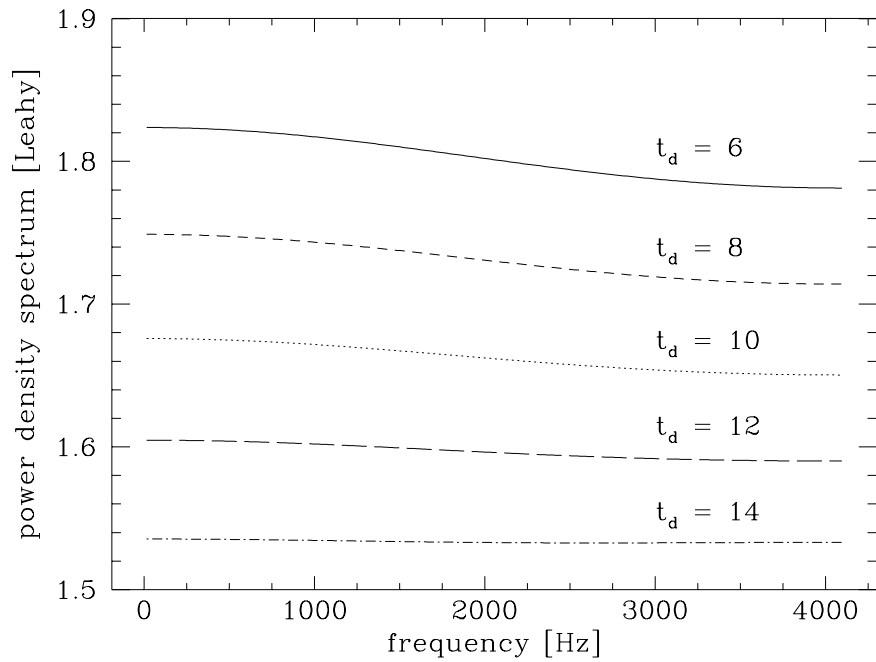

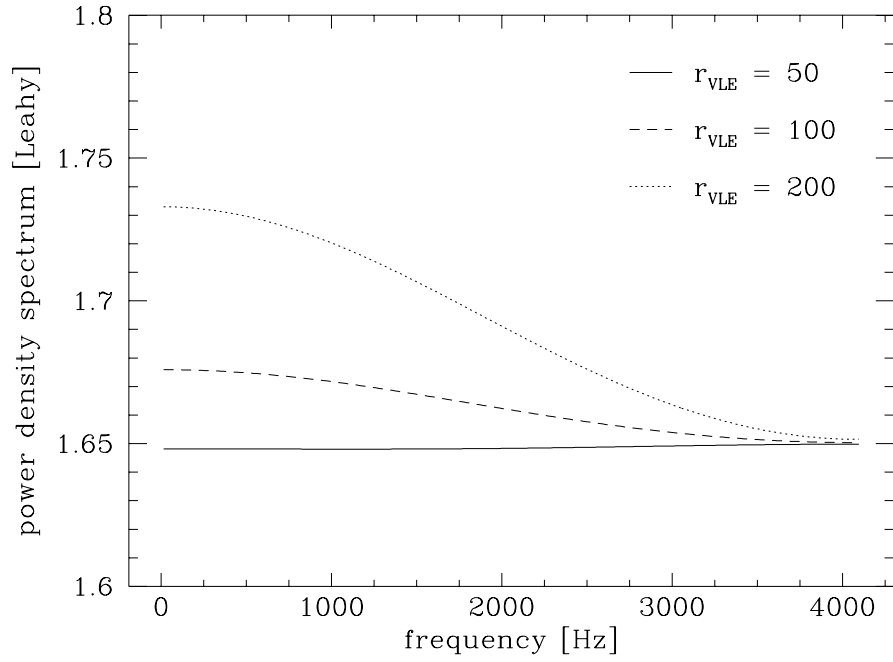Figure 4-6: Model PDS for different values of t_d [μs] and $t_b > t_d$.

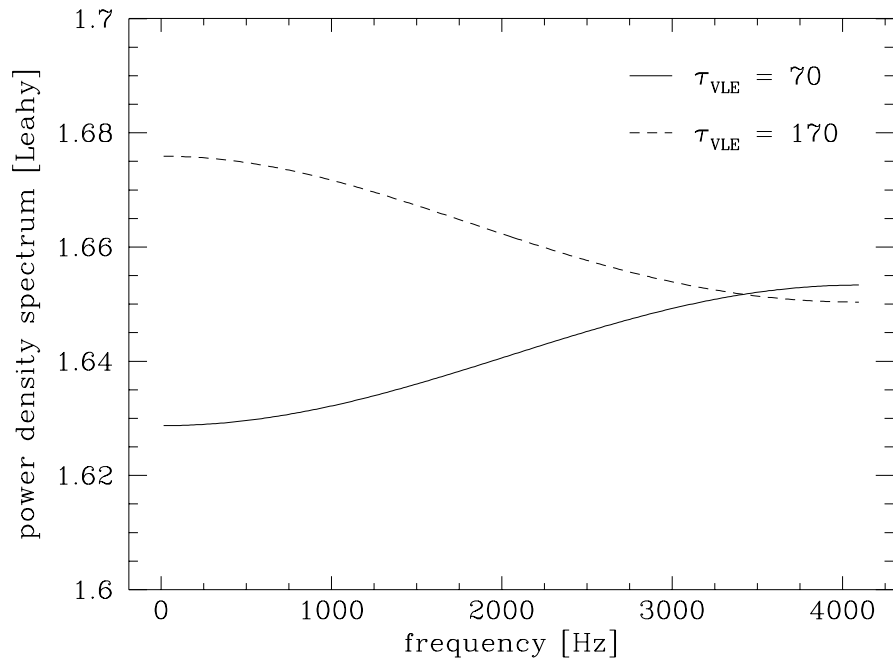Figure 4-7: Model PDS for different values of r_{VLE} [counts/sec/PCU] and $t_b > t_d$.

Figure 4-8: Model PDS for different values of τ_{VLE} [μs] and $t_b > t_d$.

Figure 4-8 plots the PDSs computed for $\tau_{\mathrm{VLE}} = 70$ μs (nominal value for VLE setting 1) and $\tau_{\mathrm{VLE}} = 170$ μs (nominal value for VLE setting 2). The other parameters are $r_o = 10000$ counts/sec/PCU, $t_d = 10$ μs, $r_{\mathrm{VLE}} = 100$ counts/sec/PCU and $r_1 = 4000$ counts/sec/PCU. By switching from VLE setting 1 to VLE setting 2, the Poisson level can switch from a rising function of frequency to a falling one. The effect of τ_{VLE}, like that of r_{VLE}, is on the order of a few hundredths in Leahy normalization.

Figures 4-9 through 4-12 feature the exact same parameter values as Figures 4-5 through 4-8, except that the time bin size t_b is now 1 μs instead of 122 μs. The first plot, Figure 4-9, shows the dependence of the Poisson level on r_o for the case $t_b < t_d$. The more complicated, sinc-like appearance of the PDS immediately distinguishes the case $t_b < t_d$ from the previous case $t_b > t_d$. Indeed, we can think of the PDSs displayed in Figures 4-5 through 4-8 as being magnified versions of the low-frequency ends of the full PDSs in Figures 4-9 through 4-12. The PDS deviates most from the zero-deadtime value of 2 in the first trough at lowest frequency. The most positive excursion is found in the first peak, which is around 80 kHz for $t_d = 10$ μs. The ripples around 2 decay thereafter as the frequency increases. In addition, there is an upward hook near zero frequency; this feature is influenced by the presence of VLEs as will be seen shortly. Figure 4-9 confirms that the observed selected rate r_o sets the amplitude, i.e. the vertical scaling, of the deviation caused by deadtime.

Figure 4-10 is the counterpart to Figure 4-6 for $t_b < t_d$. Here we see how t_d determines the horizontal scaling of the Poisson PDS. The curve for $t_d = 6$ μs appears horizontally expanded, while the curve for $t_d = 14$ μs is horizontally compressed. The deadtime effect is also larger for higher values of t_d, as it was for larger values of r_o. Thus we can conclude that the vertical scaling of the deadtime effect on the Poisson level is determined by a combination of r_o and t_b, while the horizontal scaling is largely influenced by t_d alone.
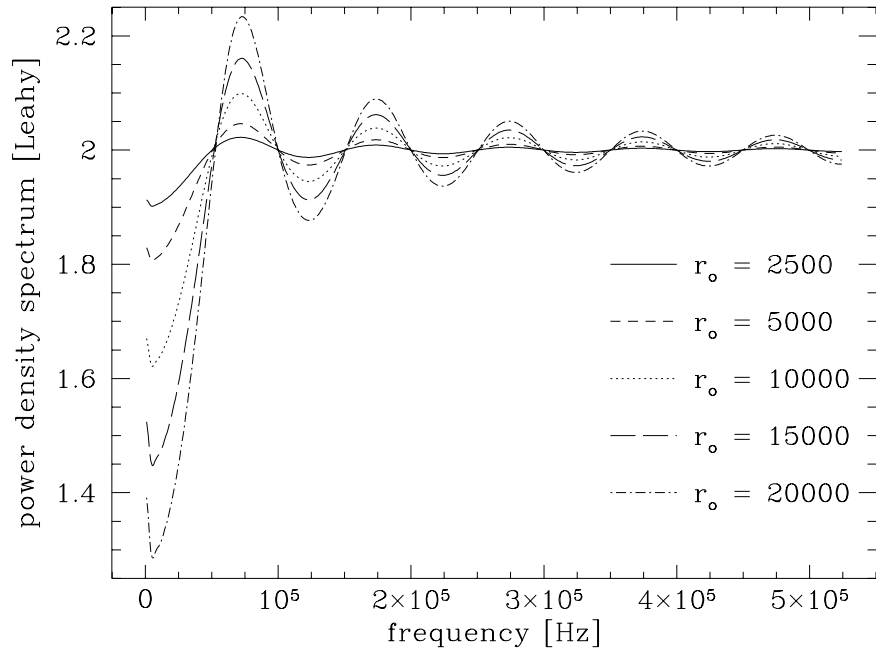
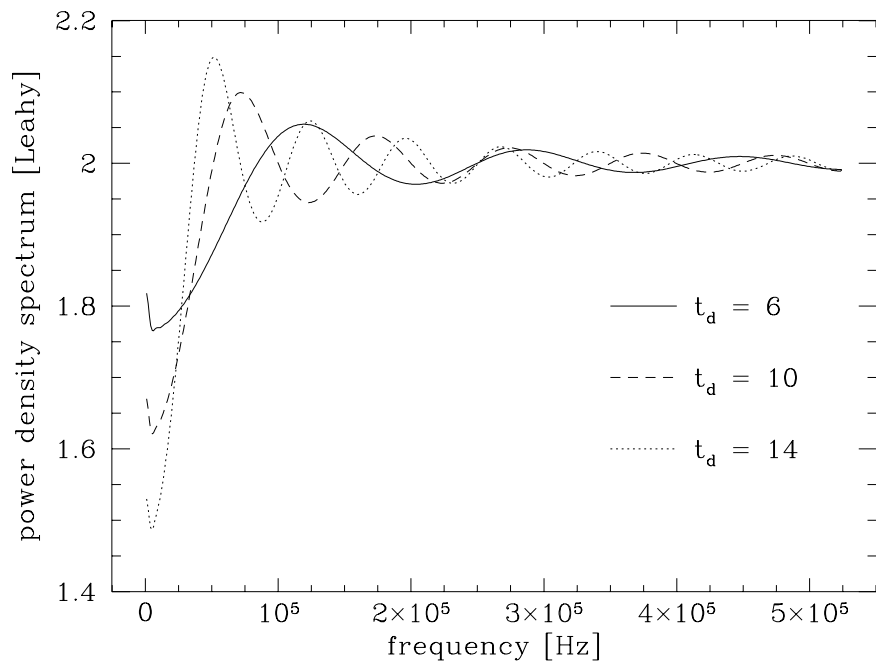Figure 4-9: Model PDS for different values of r_o [counts/sec/PCU] and $t_b < t_d$.

Figure 4-10: Model PDS for different values of t_d [μs] and $t_b < t_d$.

Figure 4-11 plots the PDSs obtained for three values of r_{VLE} and $t_b < t_d$. It is evident that the VLE effect, occurring on a timescale of τ_{VLE} that is much longer than t_d, is a low-frequency effect that does not extend beyond 10 kHz. The main difference among the three curves, although somewhat difficult to discern from the plot, is that the size of the upward hook near zero frequency is enhanced as r_{VLE} increases. This observation agrees with the r_{VLE} effect seen in Figure 4-7 when $t_b > t_d$.

Figure 4-12 shows the model PDSs corresponding to VLE settings 1 and 2 for $t_b < t_d$. As was the case in Figure 4-11, the τ_{VLE} effect is confined to the lowest frequencies. In particular, the upward hook is absent for VLE setting 1 ($\tau_{\mathrm{VLE}} = 70$ μs).

Finally, we have not included corresponding plots for the Poisson level as a function of r_1, the rate of non-selected events excluding VLEs. In actual PCA observations, the effect of r_1 on the model PDS is on the order of a few thousandths in Leahy normalization and is too small to be visually discernible. Nevertheless, r_1 does play an important role in relating the observed selected rate r_o to the inferred incident rate r_i via equation (4.14). If r_1 were absent, a key source of deadtime loss would be neglected, and the inferred incident rate would be lower than the actual incident rate. Thus, while the effect of r_1 on the Poisson level is relatively unimportant, it is still a significant contributor to deadtime.
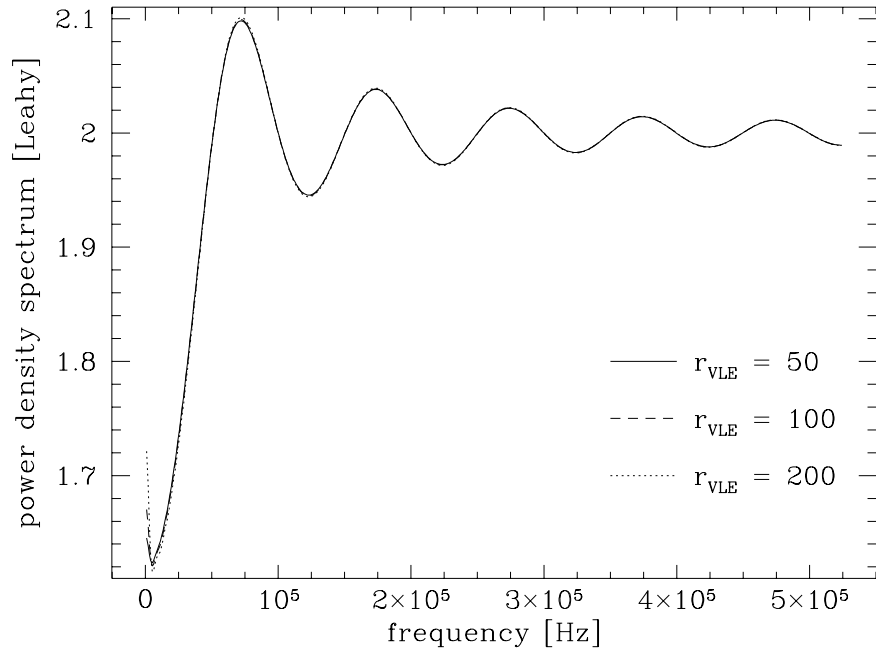
Figure 4-11: Model PDS for different values of r_{VLE} [counts/sec/PCU] and $t_b < t_d$.
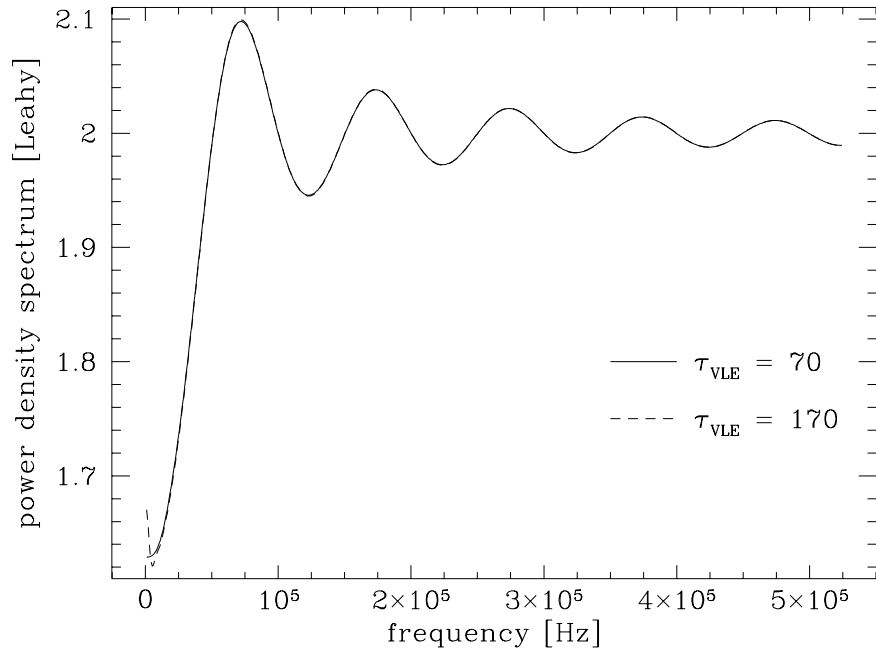

Figure 4-12: Model PDS for different values of τ_{VLE} [μs] and $t_b < t_d$.

Chapter 5

Deadtime Measurements and Model Testing

5.1 Measuring the VLE window τ_{VLE}

As discussed in Section 2.2, the RXTE PCA features a longer type of deadtime, the deadtime τ_{VLE} triggered by VLEs, in addition to the deadtime t_d associated with analog-to-digital conversion for normal events. PCA users can select one of two VLE window lengths. The first, shorter window setting is intended to decrease the deadtime impact on very bright sources, while the second, longer window is used for more modest sources and produces an output containing fewer false events. There is some disagreement concerning the exact window lengths for the two VLE settings. Zhang et al. (1996), Morgan et al. (1997), and Jahoda et al. (2006) cite values of 70 μs and 170 μs, while Jernigan et al. (2000) and Revnivtsev et al. (2000) cite 61 μs for the first setting and 142/150 μs for the second setting. In order to address these discrepancies, we performed a direct, in-flight measurement of both VLE windows using transparent mode data.

We examined nine transparent mode observations, listed in Table 5.1, of the bright X-ray source Sco X-1: four using VLE setting 1 and five using VLE setting 2. In the transparent data mode, every single event is classified and time-tagged to a precision of 1 μs (2^{-20} s). In particular, the arrival times of all VLEs are recorded. For each

VLE in the data, the time between the VLE and the next event of any kind was recorded and entered into a histogram. Due to the detector's VLE mechanism, no events should occur within τ_{VLE} after a VLE. Accordingly, the histograms should show no events below τ_{VLE} and a sudden rise at τ_{VLE}. Only a single PCU was active during these observations, allowing its deadtime to be isolated. Sco X-1 was used for the measurement of τ_{VLE} because its brightness provides ample statistics for the histograms.

VLE setting	observation ID	observation date (mm/dd/yyyy)
1	40803-01-01-00	02/19/1999
	40803-01-02-00	02/19/1999
	40803-01-03-00	02/24/1999
	40803-01-04-00	02/25/1999
2	30801-01-16-00	11/06/1998
	30801-01-21-22	11/06/1998
	30801-02-14-00	11/06/1998
	30801-03-14-00	11/07/1998
	30801-05-13-00	11/06/1998

Table 5.1: Transparent data observations of Sco X-1.

The histogram obtained for VLE setting 1 is shown in Figure 5-1 and that for VLE setting 2 is shown in Figure 5-2. From visual inspection of the rising edges in the histograms, we determine the values of the two VLE windows to be 59 μs and 137 μs, marked by dashed lines in the plots. These values are substantially lower than the nominally accepted values of 70 μs and 170 μs. They are, however, rather close to the values of 61 μs and 142 μs used by Jernigan et al. (2000).
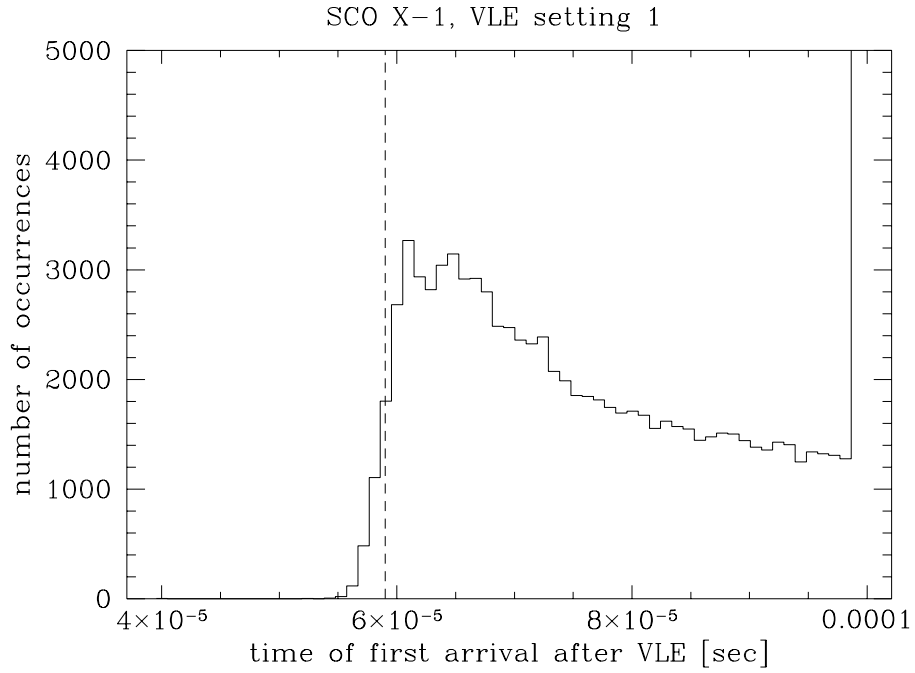
Figure 5-1: Histogram of time elapsed between a VLE and the next event for VLE setting 1. The dashed line marks $\tau_{\mathrm{VLE}} = 59~\mu$s.
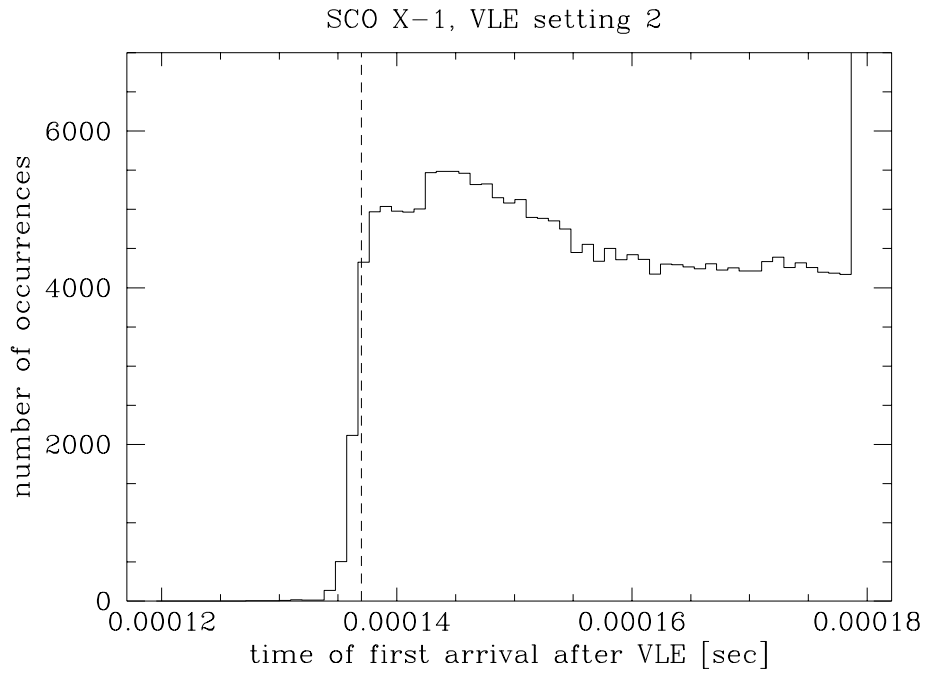

Figure 5-2: Histogram of time elapsed between a VLE and the next event for VLE setting 2. The dashed line marks $\tau_{\mathrm{VLE}} = 137~\mu$s.

5.2 Comparing deadtime models against very-high-frequency data and measuring the ADC deadtime t_d

In this section, we describe how the deadtime models were evaluated using very-high-frequency PDSs obtained from a number of astronomical sources. The fitting procedure used in testing the models also produced an accurate measurement of the ADC deadtime parameter t_d.

We used GoodXenon mode data from 4 sources: Sco X-1, GX 5-1, GX 349, and GX 9+1. A list of the observations used appears in Table 5.2. There are 30 observations for Sco X-1, 31 for GX 5-1, and 23 each for GX 349 and GX 9+1 spanning the period from April 2004 to June 2005.

source	observation IDs
Sco X-1	90024-01-nn-00
	$nn = 01$ through 31 excluding 21
GX 5-1	90024-03-nn-00
	$nn = 01$ through 29 excluding $21, 25$
	90024-03-nn-01
	$nn = 09, 13, 14, 15$
GX 349	90024-04-nn-00
	$nn = 01$ through 24 excluding 05
GX 9+1	90024-05-nn-00
	$nn = 01$ through 23

Table 5.2: GoodXenon mode observations used in testing deadtime models.

Sco X-1 was chosen for its very high flux (average observed count rates of 14000–20000 counts/sec/PCU, about 6–8 Crab equivalents). As we found in Section 4.5, the effects of deadtime on the PDS scale with the observed selected flux r_o, and many of the more subtle effects are appreciable only at high count rates. The very high flux of Sco X-1 thus presents a severe challenge for deadtime models. If a model can describe the deadtime effect for Sco X-1, then it should do even better for sources with lower, more commonplace flux levels. The latter category is represented by GX 5-1, GX 349 and GX 9+1. GX 5-1 is the second brightest persistent X-ray binary,

having count rates in the range 2000–3000 counts/sec/PCU or around 1 Crab, while GX 349 and GX 9+1 have slightly lower count rates.

GoodXenon mode offers 1 μs (2^{-20} s) timing resolution, which translates into a Nyquist frequency of 500 kHz.[1] While GoodXenon mode does allow flexible spectral analyses, its use must be limited to faint and moderately bright X-ray sources. For bright X-ray sources, it is a very expensive data mode in terms of the daily telemetry budget allowed to XTE by NASA. However, brief observations of bright sources are ideal for calibration purposes, such as testing and verifying the deadtime model of this thesis.

We computed very-high-frequency, Leahy-normalized PDSs from GoodXenon observations, containing frequencies up to 500 kHz, using the method of periodogram averaging described in Section 3.2. Specifically, a single average PDS was computed for each of the 107 observations from the four sources. We take advantage of the absence of source power for frequencies above 1–2 kHz and isolate the deadtime-modified Poisson power for fitting to deadtime models. From the quality of the fits, we determine the accuracy of the present deadtime model and compare it to previous models.

The two free parameters in the fits are the observed selected count rate r_o and the ADC deadtime t_d. The count rate r_o was varied by first reading its value from the data, r_o^{obs}, and then scaling r_o^{obs} by a correction factor r to yield $r_o^{fit} = r \cdot r_o^{obs}$. The ratio r was usually varied between 0.9 and 1.1. A ratio of $r = 1$ means that the best fit value of r_o exactly matches the observed value. A value for r close to unity indicates that the deadtime model is accurate and that all of the significant sources of deadtime are taken into account. The deadtime t_d was initially varied between 8 and 10 μs, which was soon narrowed to a smaller range centred on 8.8 μs. The best fit value for t_d is an estimate for the ADC deadtime, a very important measurement in characterizing the RXTE PCA instrument.

As for the other parameters, the VLE rate $r_{\rm VLE}$, the propane rate r_p, and the "bad"

---

[1] GoodXenon data also offers a full 256 channels of energy resolution, but we did not take advantage of this to test the deadtime model for individual energy bands beyond a preliminary stage. Please see Section 6.2 for further discussion.

rate r_b were provided by Standard 1 data accompanying the GoodXenon observations, and were fixed at their observed values throughout the fits. Either $\tau_{\mathrm{VLE}} = 59\ \mu s$ or $\tau_{\mathrm{VLE}} = 137\ \mu s$ (the values found in the previous section) were used, depending on the VLE setting in effect during the observation. The time bin size $t_b = 2^{-20}$ s and N, the number of data points included in each segment, were also taken directly from the data. The fits were evaluated using the mean squared error χ^2, defined for two degrees of freedom by

$$\chi^2 = \frac{1}{N_{pts} - 2} \sum_k \left(\frac{I_{model}[k] - I_{obs}[k]}{\sigma_I[k]} \right)^2,\qquad (5.1)$$

where $I_{obs}[k]$ is the average periodogram computed from the observations, $I_{model}[k]$ is the model for the Poisson PDS, $\sigma_I[k]$ is the statistical measurement uncertainty in $I_{obs}[k]$, and k denotes the frequency bin. The frequency spacing is linear at 1024 Hz per frequency bin. The fits included all N_{pts} frequency points above 1.5 kHz up to the Nyquist frequency. The lower frequency boundary of 1.5 kHz was chosen to avoid intrinsic power contributions from the X-ray source.

We employed three different models in fitting the data from the high-frequency PDSs. The three models were compared to determine how well they match the deadtime-modified Poisson power from real observations. Model I combines the paralyzable deadtime model of Zhang et al. (1995) with the additive VLE deadtime model of Zhang et al. (1996), as has been the practice in many RXTE research papers since their introduction (e.g. Morgan et al. 1997). The Leahy-normalized Poisson PDS for Model I is the sum of equation (27) in Zhang et al. (1995), which is the discrete-time analogue to equation (1.1) in this thesis for the case $t_b < t_d$, and equation (1.3), the additive VLE contribution due to Zhang et al. (1996). Model II differs from Model I in using the non-paralyzable deadtime model of Zhang et al. (1995) instead of the paralyzable case. The PDS for Model II may be calculated according to the procedure in Chapter 4, omitting all contributions from non-selected events (i.e. setting $r_{\mathrm{VLE}} = 0, \tau_{\mathrm{VLE}} = 0, r_1 = 0$ temporarily), and later adding equation (1.3) to the result (with r_{VLE} and τ_{VLE} now assuming their proper values). Note that r_1, which is the

rate of all non-selected events that are not VLEs, is not considered at all in Models I and II. Model III is the one developed in Chapter 4 of the thesis. It calculates the deadtime effects of selected events, VLEs, and other non-selected events in a unified manner.

We remark that a time bin size of $t_b = 1$ μs falls in the regime $t_b < t_d$, where the effect of deadtime on the Poisson PDS becomes much more complex and interesting. As seen in Section 4.5, the Poisson PDS for $t_b < t_d$ contains many waves and has approximately the shape of an inverted sinc function. The added complexity of the deadtime effect when $t_b < t_d$ poses a greater challenge to quantitative models, but at the same time instils greater confidence in a successful model.

Figures 5-3 and 5-4 depict the scatter in best fit r and t_d values for Model III, plotted as a function of the observed selected count rate r_o. A different plotting symbol is used for each source: open squares for Sco X-1, x's for GX 5-1, solid triangles for GX 349, and open circles for GX 9+1. Each point in the plots represents one observation. The dashed line in Figure 5-4 marks $t_d = 8.83$ μs. The distributions of best fit values do not show any trends with respect to the count rate r_o. The distributions tend to be wider at lower count rates because of greater statistical variation in the PDSs.

The best fit values were averaged over all of the observations for a given source, resulting in a single set of optimum values for each source and for each model. These optimum values are summarized in Table 5.3. The statistical standard deviation among the best fit parameters that are averaged is also given.
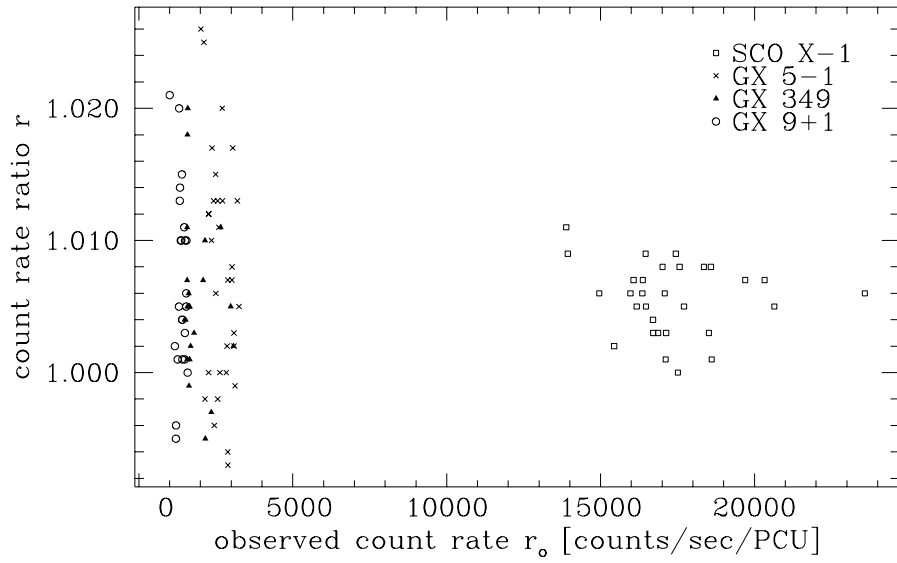
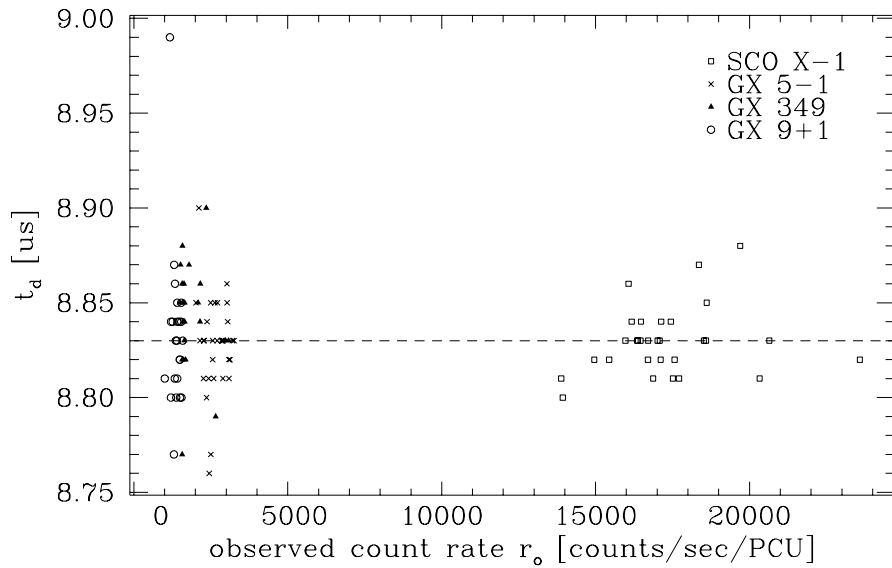Figure 5-3: Plot of optimum values for the count rate ratio r as a function of the observed count rate r_o.

Figure 5-4: Plot of optimum values for the ADC deadtime t_d as a function of the observed count rate r_o.

source	parameter	Model I	Model II	Model III
Sco X-1	r	1.013 ± 0.002	1.012 ± 0.002	1.006 ± 0.003
	t_d [μs]	8.73 ± 0.02	8.82 ± 0.02	8.83 ± 0.02
	χ^2	5.80 ± 2.36	1.08 ± 0.10	1.00 ± 0.07
GX 5-1	r	1.013 ± 0.009	1.012 ± 0.009	1.008 ± 0.008
	t_d [μs]	8.81 ± 0.02	8.83 ± 0.02	8.83 ± 0.03
	χ^2	1.02 ± 0.07	1.00 ± 0.06	1.00 ± 0.06
GX 349	r	1.010 ± 0.006	1.010 ± 0.006	1.005 ± 0.006
	t_d [μs]	8.83 ± 0.03	8.84 ± 0.03	8.84 ± 0.03
	χ^2	1.00 ± 0.08	0.99 ± 0.08	0.99 ± 0.08
GX 9+1	r	1.012 ± 0.007	1.011 ± 0.007	1.007 ± 0.007
	t_d [μs]	8.82 ± 0.04	8.83 ± 0.04	8.83 ± 0.04
	χ^2	1.00 ± 0.06	1.00 ± 0.06	1.00 ± 0.06

Table 5.3: Comparison of best fit parameters for models I, II and III.

The ratio r for Model III is about 0.005 or half a percent closer to the ideal $r = 1$ than in Models I and II. A value of r higher than 1 means that, according to the model, the recorded rate r_o^{obs} of selected events is insufficient to account for the amount of deadtime implied by the data. The lower r value for Model III indicates that a greater proportion of deadtime-causing events are being identified as compared to Models I and II.

Looking at the χ^2 values for GX 5-1, GX 349, and GX 9+1, it appears that all three models match the observed PDSs quite well. However, the χ^2 values for Sco X-1, for which the deadtime effect is much more pronounced, indicate a marked preference for non-paralyzable deadtime models (II and III). In addition, Model III shows a marginal improvement in χ^2 over Model II for Sco X-1. We conclude therefore that the deadtime t_d for the RXTE PCA is predominantly non-paralyzable, as was expected for analog-to-digital conversion.

Table 5.3 also yields the value of 8.83 μs for the ADC deadtime, which is substantially lower than the nominal 10 μs that has been assumed for the RXTE PCA (Zhang et al. 1996; Morgan et al. 1997; Jahoda et al. 2006). Our fits strongly support $t_d = 8.83$ μs as the right value to use when correcting PDSs for the ADC deadtime.

Figures 5-5, 5-6 and 5-7 show the average PDS for Sco X-1 observation 90024-01-19-00 and the best fits to the data using models I, II, and III respectively. The

data points are marked by triangles while the best fit curve is drawn as a solid line. The corresponding χ^2 values are 6.21, 1.22 and 1.09, and the improvement in the quality of the fit can be seen in the sequence of plots. In Figure 5-5, the Poisson level predicted by Model I undershoots the data by many hundredths at frequencies below 15 kHz. Model I also fails to match the peaks and troughs of the oscillations at higher frequencies. In Figure 5-6, Model II improves significantly upon Model I, although there is still some minor undershoot at the lowest frequencies and the first peak. A non-paralyzable deadtime model is clearly a better fit to the Sco X-1 data than a paralyzable model. Model III in Figure 5-7 is a marginal improvement over Model II when comparing the lowest frequencies and the first peak. However, the best fit values for r are 1.010 ± 0.003 for Model II and 1.003 ± 0.003 for Model III, so Model III is also better because its r value is closer to unity.

Figure 5-8 is a comparison of the Poisson PDSs produced by Models I, II and III given an identical set of input parameters. The input parameters are $r_o = 17000$ counts/sec/PCU, $t_d = 8.83$ μs, $r_{\mathrm{VLE}} = 100$ counts/sec/PCU, $\tau_{\mathrm{VLE}} = 137$ μs, $r_1 = 6800$ counts/sec/PCU, $t_b = 2^{-20}$ s and $N = 2^{20} = 1048576$. As in Figure 5-5, the PDS of Model I can be distinguished from those of Models II and III by the undershoot at low frequencies and the smaller amplitude of the oscillations at higher frequencies. The difference between Models II and III is very subtle; the PDS of Model II appears to be a slightly scaled-down version of that of Model III. This difference makes sense, since the value of r, and hence r_o, for Model II had to be slightly higher than that of Model III when fitting to the same set of data.

Figure 5-5: Average PDS from Sco X-1 observation 90024-01-19-00 (triangles) and best fit using model I (solid line).
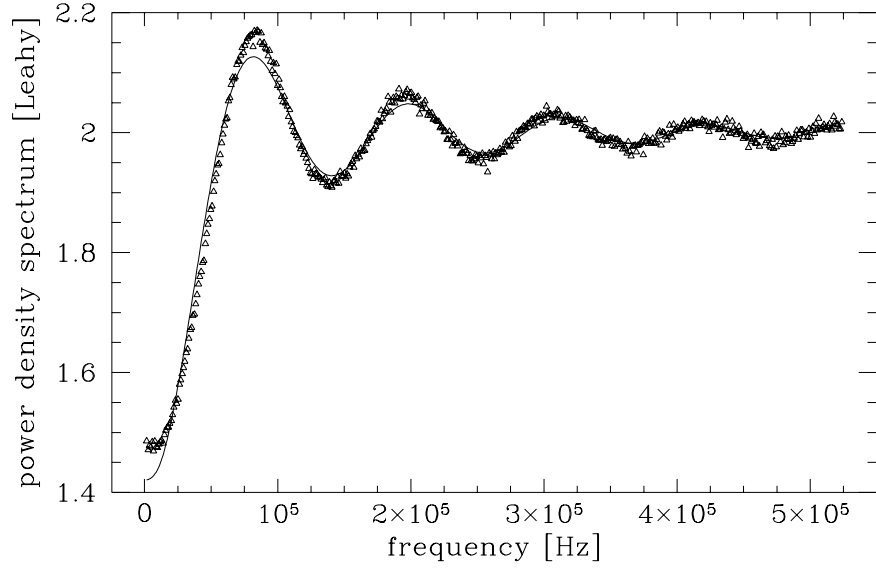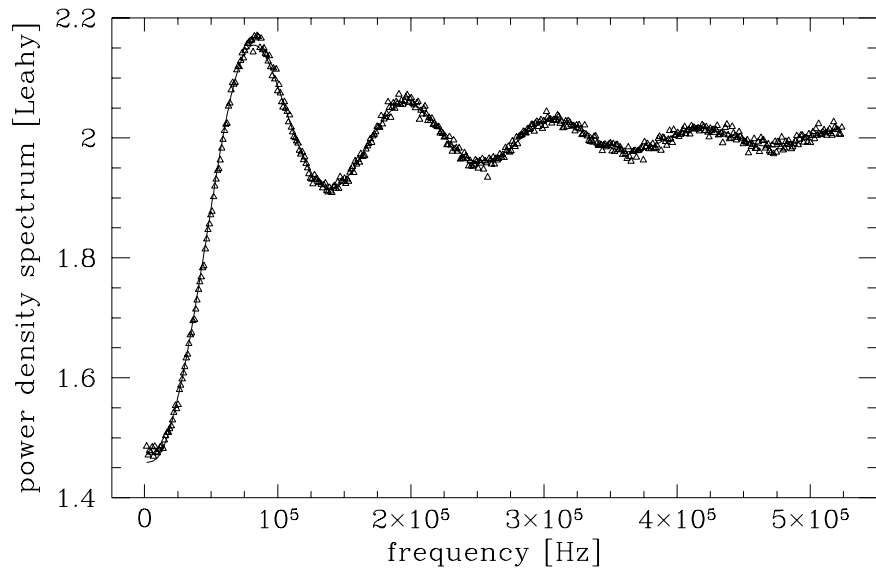

Figure 5-6: Average PDS from Sco X-1 observation 90024-01-19-00 (triangles) and best fit using model II (solid line).
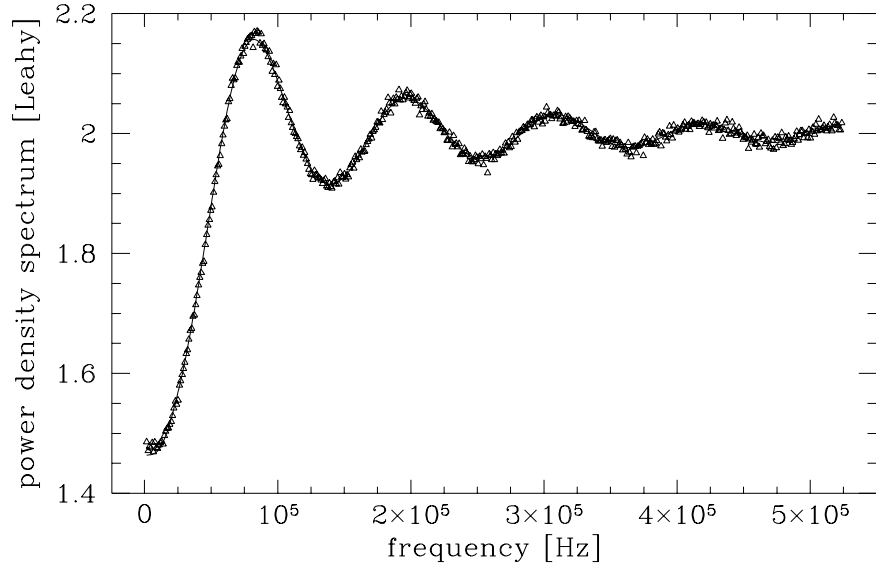
Figure 5-7: Average PDS from Sco X-1 observation 90024-01-19-00 (triangles) and best fit using model III (solid line).
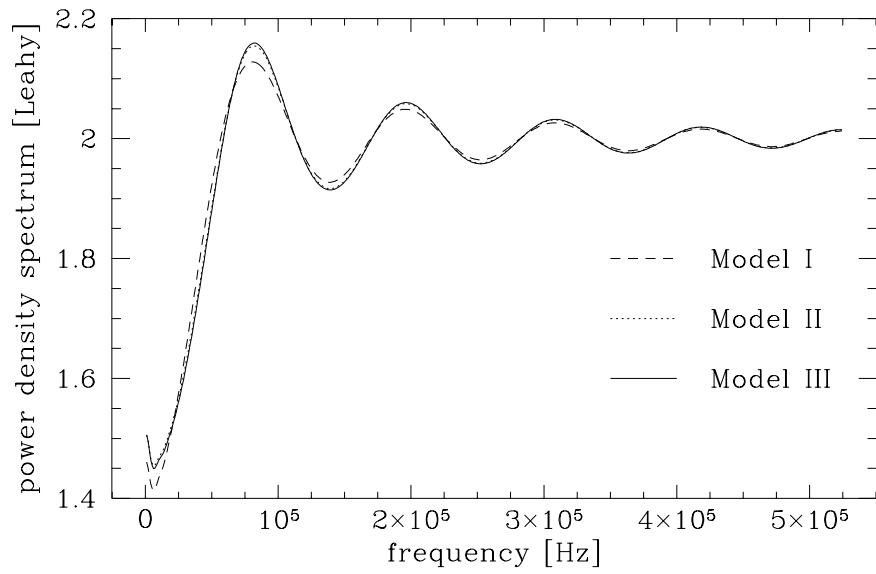

Figure 5-8: Comparison of Poisson PDSs produced by Models I, II and III given identical input parameters.

Chapter 6

Additional Deadtime Effects for Future Study

6.1 Non-stationarity of highly variable sources

In deriving the deadtime model in this thesis, we assumed that the incident events follow a stationary Poisson process characterized by a constant rate, which is a standard approach in constructing such models (Zhang et al. 1995; Vikhlinin et al. 1994). The assumption of stationarity works well when the amount of source variability is modest and the source flux is approximately constant. However, astrophysical sources are not always well behaved and can exhibit considerable fluctuations relative to their average flux. The deadtime effect for a highly variable source may be very different from the deadtime effect for a pure Poisson process, even though the average flux is the same in both cases. For highly variable sources, we should expect some difficulties in using a model based on a stationary Poisson process since the model does not match the source's behaviour.

A well-known example of an astrophysical source capable of extreme variability is the BHB GRS 1915+105. We analyzed a small collection of RXTE PCA observations of GRS 1915+105 in the "θ state" taken in September and October 1997. According to the classification scheme of Belloni et al. (2000) for the types of variations of GRS 1915+105, the θ state is characterized by an irregular pattern of build-ups and sudden

releases in the source flux. A representative light curve depicting the dramatic flux variations in the θ state is shown in Figure 6-1. The count rate ranges from a minimum near 8000 counts/sec to a maximum over 40,000 counts/sec. In comparison, the light curve for GX 5-1 in Figure 6-2 exhibits much smaller relative variations.

The observations of GRS 1915+105 in the θ state were taken at a time resolution of approximately 500 μs (2^{-11} s). No energy selection was performed on the events; all events from the full observable energy spectrum were included. We employed five different methods of calculating PDS estimates (average periodograms) from the GRS 1915+105 θ state timing data:

1. segment length of 256 seconds, periodogram averaging over an observation, followed by subtraction of deadtime-corrected Poisson level based on count rates averaged over the observation

2. same as #1, but with a segment length of 64 seconds

3. same as #1, but with a segment length of 16 seconds

4. same as #1, but with a segment length of 1 second

5. segment length of 1 second, subtraction of Poisson level based on count rates *within each segment*, followed by averaging corrected periodograms over an observation

The Leahy-normalized PDSs obtained through Methods 1–3 are presented in Figure 6-3. Method 1 is represented by a solid line, Method 2 by a dotted line, and Method 3 by a dashed line. The PDSs corresponding to Methods 4 and 5 are shown in Figure 6-4. Method 4 is represented by a solid line and Method 5 by a dotted line. The PDSs are plotted on a linear scale and the limits have been chosen in order to focus on the high-frequency residual power.

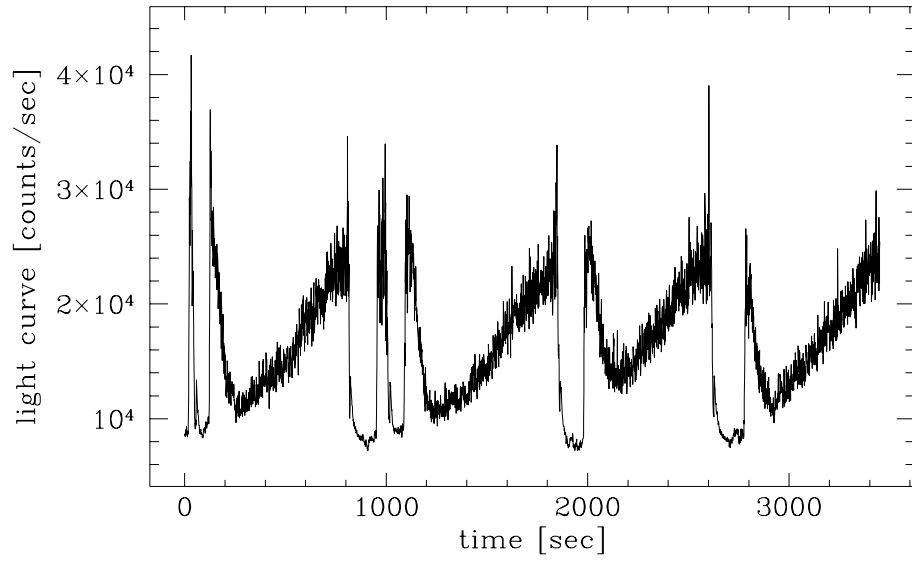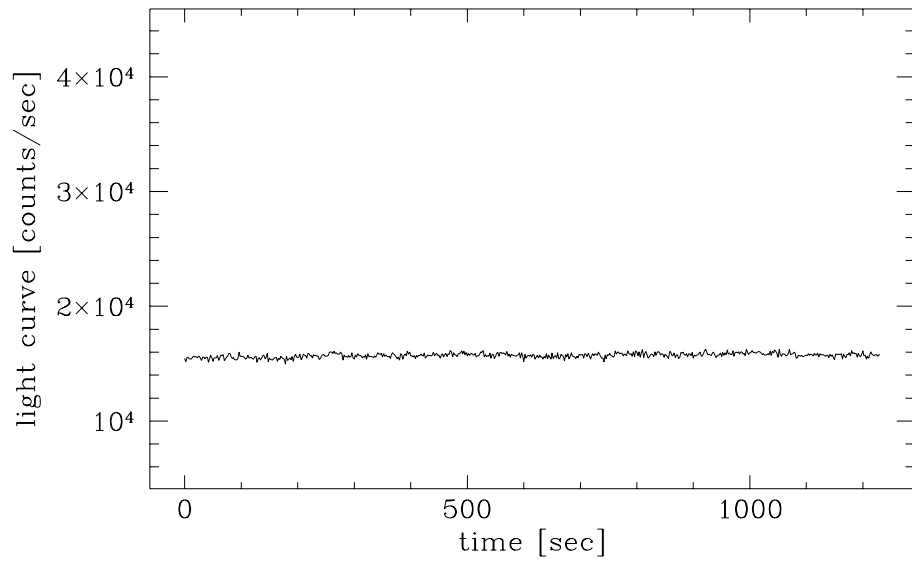Figure 6-1: Sample light curve for GRS 1915+105 in the θ state.

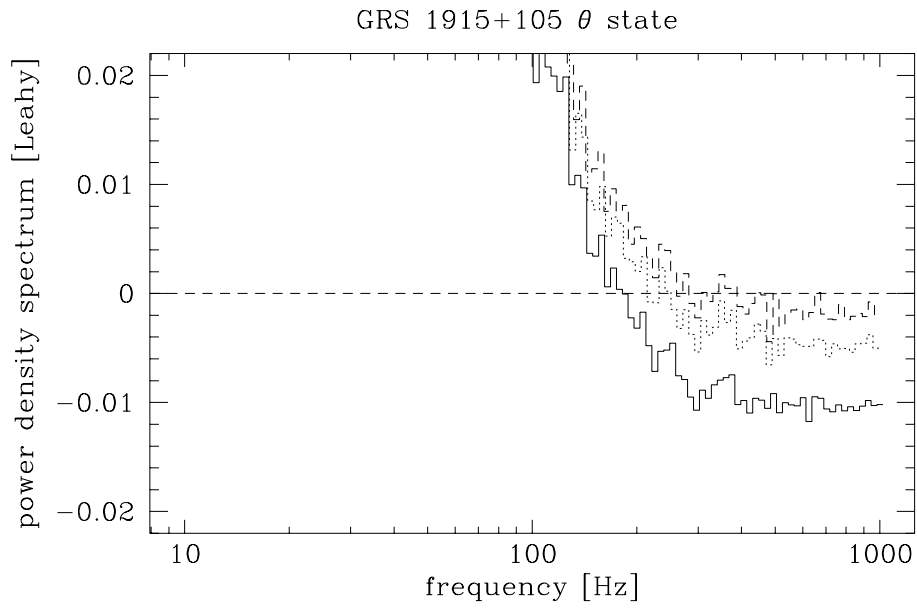Figure 6-2: Sample light curve for GX 5-1.

Figure 6-3: Leahy-normalized average PDS for GRS 1915+105 in the θ state using Methods 1 (solid line), 2 (dotted line), and 3 (dashed line).
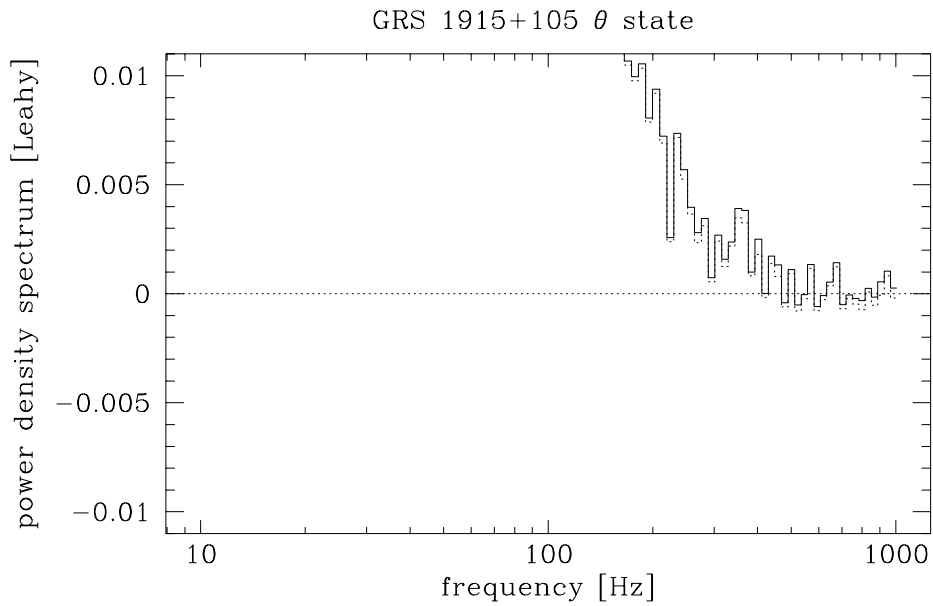

Figure 6-4: Leahy-normalized average PDS for GRS 1915+105 in the θ state using Methods 4 (solid line) and 5 (dotted line).

In all cases, the PDS decays to an approximately constant level for frequencies around 400 Hz and higher. The value of the constant level steadily decreases as the segment length increases, going from near zero for 1-second segments to -0.01 for 256-second segments. Since the intrinsic source power must be positive at all frequencies, we conclude that the negative residuals in Figure 6-3 are caused by over-subtraction of the Poisson power. In other words, the deadtime model overestimates the correction to the Poisson power for GRS 1915+105 and the extent of the overestimate clearly increases with the length of the transformed segments.

The dependence of the negative residual on segment length suggests that the culprit could be the amount of source variability observed within a segment. As the light curve in Figure 6-1 indicates, if an observation is divided into longer segments (e.g. 256 seconds), the full extent of the variability of GRS 1915+105 is more apparent within each segment. Segments that, in the extreme case, contain one of the sharp drops in flux can hardly be called stationary. We expect that the deadtime model should fail when applied to long segments containing clearly non-stationary behaviour. In contrast, dividing an observation into short segments (e.g. 1 second) hides much of the long-term, large-amplitude variation if each segment is viewed in isolation. The deadtime model should perform better on short, quasi-stationary segments. In summary, the longer the segment length, the more source variability is captured within each segment, and the greater the discrepancy between the deadtime model and the actual Poisson power.

We can show through the following semi-quantitative analysis that the residual power should be correlated with a particular measure of source variance. We focus our attention on the dependence of the Poisson power on the observed selected count rate $r_o(t)$, which is no longer constant in time. Furthermore, we ignore all of the frequency-dependent terms in equation (4.17) (all of the $m \neq 0$ terms in the summation), and approximate the Poisson power by the constant $m = 0$ term over all frequencies. Given these simplifications, the deadtime-corrected Poisson power in Leahy units is an approximately linear function of r_o (for modest values of r_o):

$$P_{Poisson}^{Leahy} \approx 2(1 - \alpha r_o), \tag{6.1}$$

where $\alpha > 0$ is a parameter for the slope. Note that $P_{Poisson}^{Leahy} \to 2$ as $r_o \to 0$ in the absence of deadtime. The Poisson power in raw, un-normalized units is a quadratic function of r_o, obtained from (6.1) by multiplying by $r_o/2$:

$$P_{Poisson}^{raw} \approx r_o - \alpha r_o^2. \tag{6.2}$$

The observed count rate $r_o(t)$ may vary significantly over the course of a data segment. As an intuitive guess, we suppose that the resulting Poisson power in un-normalized units is the time average of equation (6.2) over the segment:

$$\langle P_{Poisson}^{raw} \rangle_S \approx \langle r_o \rangle_S - \alpha \langle r_o^2 \rangle_S, \tag{6.3}$$

where the subscript S denotes an average over a segment. In processing data from observations, the raw PDSs from individual segments are Leahy-normalized (multiplied by $2/\langle r_o \rangle_S$) before being averaged over an observation (periodogram averaging):

$$\left\langle \left\langle P_{Poisson}^{Leahy} \right\rangle_S \right\rangle_O \approx 2 \left(1 - \alpha \left\langle \frac{\langle r_o^2 \rangle_S}{\langle r_o \rangle_S} \right\rangle_O \right), \tag{6.4}$$

where the subscript O denotes an average over an observation.

On the other hand, our model for the Poisson power uses the count rate averaged over an observation in equation (6.1):

$$P_{Poisson}^{model} = 2 \left(1 - \alpha \langle \langle r_o \rangle_S \rangle_O \right). \tag{6.5}$$

Subtracting equation (6.5) from (6.4) yields the residual power $\Delta P_{Poisson}^{Leahy}$:

$$\begin{aligned}
\Delta P_{Poisson}^{Leahy} &\approx -2\alpha \left\langle \frac{\langle r_o^2 \rangle_S - \langle r_o \rangle_S^2}{\langle r_o \rangle_S} \right\rangle_O, \\
&\approx -2\alpha \left\langle \frac{var_S(r_o)}{\langle r_o \rangle_S} \right\rangle_O,
\end{aligned} \tag{6.6}$$

where we have identified $var_S(r_o)$ as the variance of $r_o(t)$ computed for each segment. Thus, the foregoing analysis predicts a linear relationship between the residual power and the *average source variance* defined by $\langle var_S(r_o)/\langle r_o \rangle_S \rangle_O$.

The average source variance $\langle var_S(r_o)/\langle r_o \rangle_S \rangle_O$ can be estimated from discrete-time count data. Equivalently, it can be estimated in the frequency domain from the periodogram corresponding to the same segment. Specifically, the average source variance is approximately equal to the "integral" over all positive frequency bins ($k = 1$ to $k = N/2$) of the Leahy-normalized average periodogram $\langle I_{source}^{Leahy}[k] \rangle_O$:

$$\left\langle \frac{var_S(r_o)}{\langle r_o \rangle_S} \right\rangle_O \approx \frac{1}{Nt_b} \sum_{k=1}^{N/2} \langle I_{source}^{Leahy}[k] \rangle_O . \qquad (6.7)$$

The factor $1/Nt_b$ is the width of a frequency bin, hence the interpretation of the sum on the right-hand side as an approximation to an integral. In addition, only the source component of the average periodogram is integrated, as indicated by the subscript "source", since the average source variance does not include the variance due to Poisson statistics. The source component of the average periodogram can be obtained by subtracting the Poisson power from the total average periodogram.

Using equation (6.7), the average source variances corresponding to Methods 1–5 were computed by integrating their Poisson-subtracted, Leahy-normalized average periodograms over all positive frequencies. In doing so, the error due to the original over-subtraction of the Poisson level was ignored as it has only a small effect on the integrals. The high-frequency residual power (Leahy normalization, averaged above 500 Hz) is plotted against the average source variance in Figure 6-5. The correlation between the two quantities is quite clear. The predicted linear relationship is drawn as a dashed line in Figure 6-5. The dashed line has slope $-2\alpha \approx -7.5 \times 10^{-6}$ corresponding to the average flux of GRS 1915+105 and is offset so that it passes through the leftmost data point. Except for the rightmost data point, there is rough agreement between the data and the prediction despite the crudeness of the analysis.

We can also understand why Methods 4 and 5, which differ only in the order of Poisson subtraction and periodogram averaging, give comparable residuals. In

81
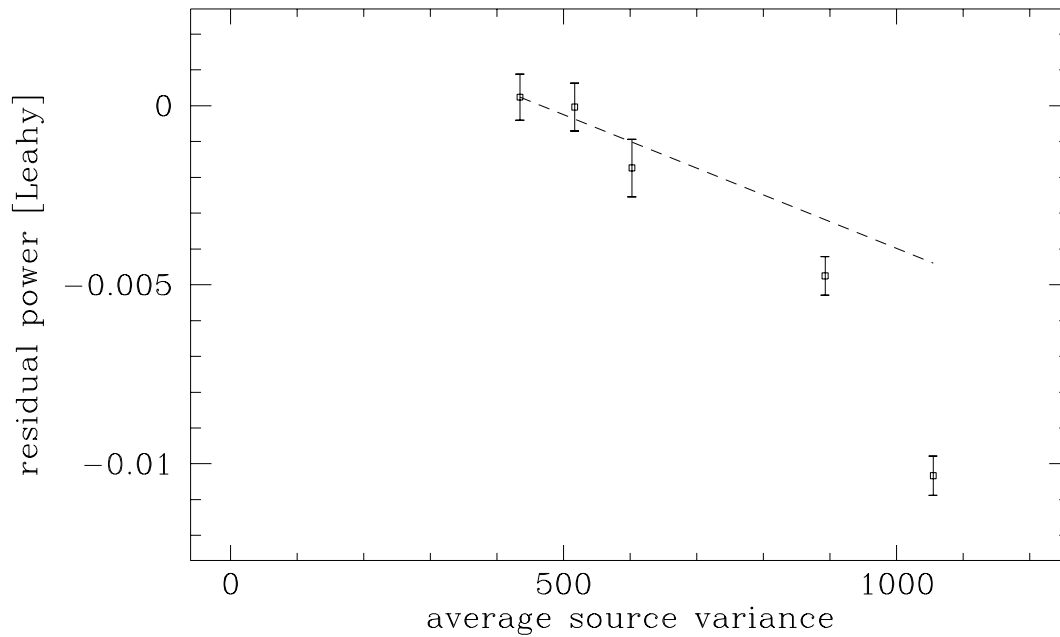
Figure 6-5: High-frequency residual power as a function of average source variance for GRS 1915+105 in the θ state.

Figure 6-5, Methods 4 and 5 are represented by the leftmost and second-leftmost data points respectively and have similar average source variances. Since both methods use a segment length of 1 second, they should theoretically experience the same average source variance. The slight discrepancy between the two methods is due to the different ways in which the Poisson component is subtracted from the total PDS. In Method 4, the Poisson subtraction is performed at the end according to rates averaged over the observation. In Method 5, the Poisson subtraction is performed on each segment according to rates within the segment.

6.2 Dependence of deadtime model on energy of events

In this thesis, we have undertaken only a preliminary exploration of the dependence of detector deadtime on the energies carried by events. The measurement of dead-time parameters and the testing of the model in Chapter 5 were done for the full energy spectrum (the sum band) permitted by the PCA. However, it is very common practice in X-ray astronomy to select only events within a certain energy band for further analysis. It is doubtful that a deadtime model with parameters adapted to the sum band will continue to be accurate for an individual band. In order to fully resolve this question, it will be necessary to systematically test the deadtime model for different energy bands in isolation. Significant discrepancies among the individual bands should then be incorporated into the deadtime model.

We should note that the current model, despite possible inaccuracies, is equipped to handle the case of events selected from a single energy band. Events in other energy bands are treated as non-selected events, equivalent for instance to propane events, since they also compete with events in the selected energy band for access to the ADC.

It is known that the detector gain, the ratio of the collected charge to the incident event energy, suffers small temporary drops immediately after every incident event. We recall from Section 2.1 that each incident event causes an avalanche of electrons to be drawn toward an anode. At the same time, the much heavier and much slower positive ions are left behind and must be collected by the cathode. Before the ions are absorbed by the cathode, they can attract electrons from a subsequent event and prevent them from reaching an anode. Thus, during the time it takes for ions to reach the cathode following an event, i.e. the ion drift time, the number of electrons collected by the anodes is reduced and the gain appears to decrease.

The gain decreases could cause a rapidly following event to yield a lower than normal pulse height and be misinterpreted as having a lower energy, and thus can lead to differences in the PDSs of different energy bands. In particular, a lower than

expected level of high-frequency variability would be found for the higher energy bands, since the effect requires that the events be separated by short time intervals, and the demotion is always toward lower energy. For example, the lowest energy band could experience a surplus of high-frequency power because of the events gained from higher energy bands, while the higher energy bands would experience a high-frequency power deficit because of the events lost. The timescale of the drops in gain is related to the ion drift time and will need to be measured or inferred. It may then be possible to model the amount of high-frequency power lost from higher energy bands and gained by lower energy bands. We note that this effect can be especially important for sources with steep energy spectra where the count rate is high at the boundaries between energy channels.

As motivating evidence for the effect of temporary drops in detector gain, we present in Figures 6-6 and 6-7 two PDSs taken from observations of GRS 1915+105 in the θ state. Figure 6-6 shows the PDS for events in the 2–6 keV band, the lowest energy band in this analysis, while Figure 6-7 shows the PDS for events in the next highest 6–10 keV band. Both PDSs were calculated using Method 4, i.e. periodograms from 1-second data segments were averaged and then corrected for deadtime and Poisson statistics. Using 1-second segments avoids most of the problems associated with high source variability, as demonstrated in Section 6.1. Figure 6-6 clearly indicates positive, high-frequency residual power for the 2–6 keV band. Conversely, the PDS in Figure 6-7 for the 6–10 keV band contains negative residual power at high frequencies. These plots support the theory that short-term decreases in detector gain reduce the high-frequency variability in higher energy bands while increasing the same in lower energy bands by a similar amount. Nevertheless, more sources need to be tested to see if a similar pattern holds.
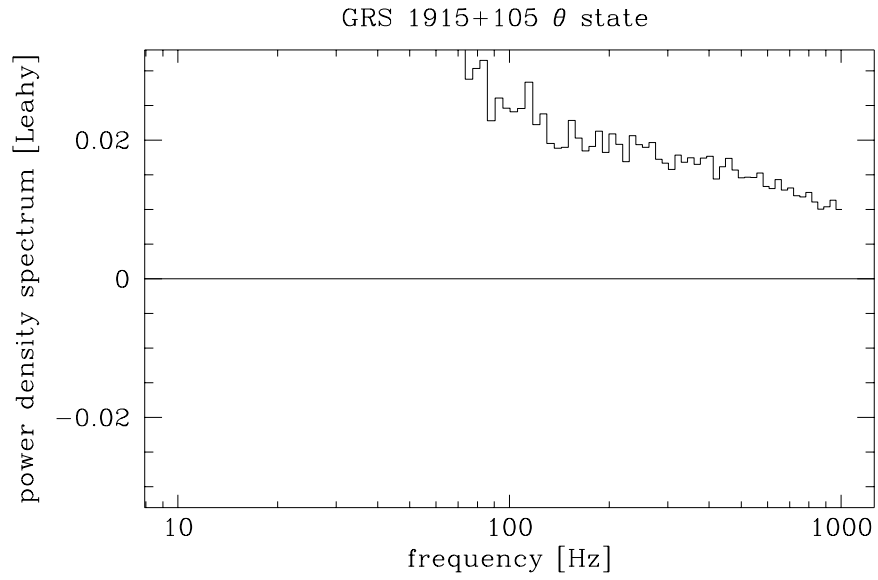
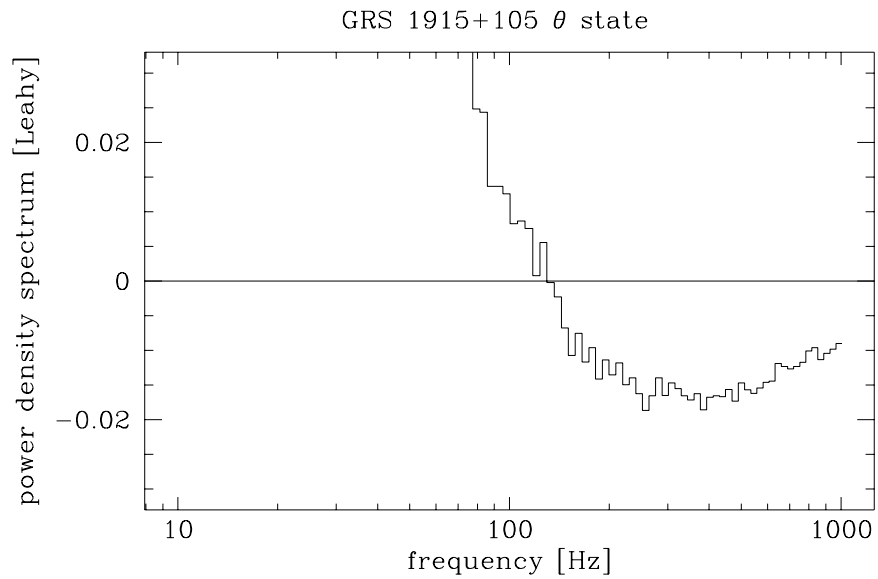Figure 6-6: Leahy-normalized average PDS for the 2–6 keV band of GRS 1915+105 in the θ state

Figure 6-7: Leahy-normalized average PDS for the 6–10 keV band of GRS 1915+105 in the θ state.

One possible approach to accounting for the differences among individual energy bands is to determine effective r and t_d values for each band. As a reminder, the correction factor r is the ratio between the observed count rate r_o and the count rate that is actually input into the deadtime model. While the values $r = 1.007$ and $t_d = 8.83\,\mu$s work well for the sum band, proper deadtime correction for the individual bands could require slightly different values. For example, some preliminary tests involving energy bands higher than 10 keV indicate that a value of $t_d = 9.0$–9.1 μs could be more appropriate. In addition, the observed count rate generally decreases with increasing energy, so that deadtime corrections become less important for data from higher energy bands.

6.3   Multiple photon coincidences

As described in Section 2.2, the RXTE PCA features an anti-coincidence mechanism that eliminates events deemed to be too close together. Specifically, if events A and B occur on different measurement chains, and if event B follows event A by less than the anti-coincidence window $t_a \approx 3$–4 μs, both events are declared "bad" and are excluded from the selected events. The anti-coincidence feature is designed to rule out particles and other non-photon events that tend to generate multiple ionizations in quick succession. However, if the source is very bright, there is a non-negligible probability that two photons will arrive within t_a of each other, causing the PCA to mistakenly disqualify both photons. The model of this thesis does not distinguish multiple photon coincidences from legitimate non-photon events, even though photon coincidences should be admitted as part of the source flux. We should point out that if the anti-coincidence feature did not exist, the second and third events of a multiple photon coincidence would still be eliminated because they arrive while the ADC is busy with the first event ($t_d > 2t_a$). The action of the anti-coincidence monitor eliminates the first event as well.

The anti-coincidence mechanism is expected to have a secondary impact on the Poisson PDS. Intuitively speaking, anti-coincidence targets regions of high-frequency

variability where the source photons tend to cluster together. The effect should scale with the count rate squared since it depends on the rate of multiple photon coincidences. A quantitative model for the anti-coincidence effect could be helpful for very bright sources such as Sco X-1. As a first step, approximate calculations of multiple coincidence rates can be found in Jahoda et al. (1999).

6.4 Ordering of deadtime processes

As explained in the section on the deadtime model, the self-deadtime among selected events is thought of as occurring before the interference effects due to VLEs, propane, and other non-selected events. However, the opposite order is closer to how events are processed in the actual detector. The order assumed in the model was chosen mainly for analytical convenience, and because the discrepancy should be negligible for all but the highest count rates.

We outline here the steps needed to analyze the more realistic ordering of deadtime processes. First, the interference by non-selected events would remove some events from the stream of incoming selected events, which is again assumed to be Poisson. The probability of detecting a selected event after undergoing interference will carry a slight dependence on the time since the last selected event, unlike a Poisson process where no such memory exists. The reason for the slight memory effect can be understood intuitively by noting that the detection of a selected event guarantees in the short term the absence of masking by non-selected events. The resulting stream of selected events would no longer follow a Poisson process. One would have to repeat the calculations of Zhang et al. (1995) to determine the effect of self-deadtime among the surviving selected events, but with a non-Poisson process instead, a derivation that could be quite difficult. Once the continuous-time autocorrelation function for the surviving selected events is determined, the discrete-time PDS could be calculated via equations (3.16) and (3.15) as before.

Chapter 7

Preliminary Comparison of Power Density Spectra from Black Hole Binaries and Neutron Star Binaries

In Section 1.3, we described a method proposed by Sunyaev & Revnivtsev (2000) for distinguishing BHBs from NSBs on the basis of their PDSs. Specifically, Sunyaev & Revnivtsev observed that PDSs of BHBs do not contain any significant power beyond several hundred Hz, whereas the PDSs of NSBs often extend into the kHz range. They claim therefore that BHBs and NSBs may be distinguished by the high-frequency cutoffs of their intrinsic power.

As we noted in Sections 1.3 and 1.4, there are two concerns with their approach. First, Sunyaev & Revnivtsev selected BHB observations from the hard X-ray state only (see Section 1.2 for description of X-ray states). Since there are large quantitative and qualitative differences among PDSs from different states (Remillard & McClintock 2006; McClintock & Remillard 2006), a more complete study must consider all of the states of both BHBs and NSBs. The omission of other states appears worse when one recalls from Section 1.2 that the BHB hard state is now clearly associated with the presence of a quasi-steady radio jet (Fender 2006; Remillard & McClintock 2006). The jet is thought to distort the radiation geometry of the inner accretion disk where high-frequency variability originates. Therefore, we might expect diminished

high-frequency continuum power in the hard state as compared with other states.

The second concern involves the deadtime modelling of Sunyaev & Revnivtsev (2000). In subtracting the Poisson component of the PDS, they fitted a paralyzable deadtime model, specifically Model I of Section 5.2, to the high-frequency portion of the total PDS. As shown in Section 5.2, Model I is less accurate than the model developed in this thesis. More importantly, the fitting process appears to be problematic because it pre-supposes a cut-off frequency and attempts to minimize the power beyond the cut-off to the extent that the model allows. The cut-off frequency chosen for NSBs is higher than that for BHBs (600 Hz versus 50 Hz), potentially biasing the results.

In this chapter, we report preliminary findings that address the two issues of X-ray state selection and deadtime modelling. The comparisons between BHBs and NSBs are re-examined in light of these refinements. We computed average PDSs for a number of BHBs and NSBs. A single PDS was obtained for each observation by periodogram averaging and Poisson subtraction using the new deadtime model and without any fitting. The BHB observations were classified by X-ray state and luminosity relative to their X-ray state. For each BHB, all of the PDSs belonging to a given X-ray state and luminosity were averaged. The NSBs chosen belong to the high-luminosity class known as "Z sources", which are named for the shape of the tracks they exhibit in plots of hardness colour versus luminosity (van der Klis 2006). For each Z source, we averaged PDSs from each of the three branches of the "Z", called the horizontal branch, normal branch, and flaring branch, where this list is given in order of increasing mass accretion rate, under current interpretations.

We should comment on the difficulty of establishing cut-off frequencies in PDSs by quantitative means. As we noted in Section 1.3, cut-off frequencies in source PDSs are strongly sensitive to the deadtime model used for Poisson subtraction. If the deadtime model is perfect, the power should be essentially zero above the cut-off frequency, with positive and negative fluctuations about zero within the statistical uncertainty. However, small systematic errors in the deadtime model can shift the level of true zero power and complicate the determination of cut-off frequencies. At the required level

of accuracy, even the improved deadtime model is vulnerable to systematic errors. A number of known sources of systematic error were discussed in Chapter 6. Of these known effects, the error due to the non-stationarity of astronomical X-ray sources is especially important since it affects nearly all sources to some degree. In addition, the model inevitably does not take into account other as-yet-unidentified effects, e.g. non-ideal behaviour in the gas chamber or analog electronics.

True zero power may also be obscured by finite-length effects in Fourier transform analysis. The finite length of transformed data segments leads to smearing of features in the PDS. In particular, if the true PDS falls to zero beyond a cut-off frequency, the observed PDS will decay as f^{-2} instead. In principle, cut-off frequencies could be detected by plotting the PDS on a logarithmic scale and searching for breakpoints where the slope becomes -2. However, except for the hard state, regions with -2 slope are not particularly common in practice and may not extend over a wide enough frequency range to be conclusive. In recognition of the difficulties of quantitative analysis, we follow Sunyaev & Revnivtsev in determining high-frequency cut-offs graphically.

We begin by comparing PDSs from the three BHB states: thermal, hard, and steep power-law (SPL). In the left panel of Figure 7-1, we present PDSs for the BHB XTE J1550-564 in all three states. In the right panel, we show PDSs for the BHB GRO J1655-40 in the thermal and SPL states, and for the BHB GX 339-4 in the hard state. In both panels, the thermal state is represented by a dashed line, the hard state by a dotted line, and the SPL state by a solid line. RMS normalization is used throughout this chapter.
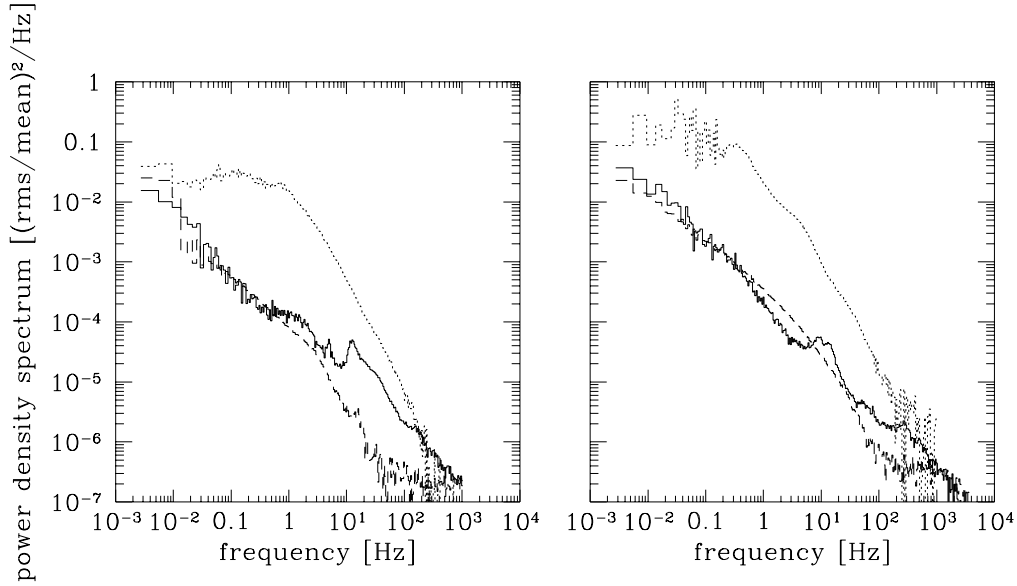
Figure 7-1: Comparison of PDSs from the three BHB X-ray states for XTE J1550-564 (left panel), GRO J1655-40 (right panel, thermal and SPL states) and GX 339-4 (right, hard state). The X-ray states are represented by the following line types: thermal (dashed), hard (dotted), and SPL (solid).

The hard state is distinguished as having the highest power at low frequencies. Beginning at several Hz in the PDS of XTE J1550-564, the continuum power is seen to decay with a logarithmic slope near -2. The PDS of GX 339-4 also has a slope close to -2 from ~ 50 Hz to 200 Hz, and contains gaps above 200 Hz that indicate negative fluctuations in the source power. These observations suggest a relatively low cut-off to the substantial low-frequency power of the hard state.

The PDSs for the thermal state are smoothly decaying for the most part. In contrast, the PDSs for the SPL state are more varied and change slope often. They contain a number of QPOs, both at lower frequencies in the vicinity of 10 Hz, and at higher frequencies in the tens and hundreds of Hz. In general, the SPL state has the strongest high-frequency content in the range 10 Hz to 1 kHz, assuming that the -2

92

slopes of the hard state are finite-length artifacts. It is difficult to determine whether or not the continuum power extends beyond 1 kHz in the SPL state as the amplitudes become comparable to the statistical uncertainty.

The PDSs of Z sources in different branches also appear to differ in systematic ways, as seen in Figure 7-2. The left panel features the Z source GX 17+2 while the right panel features GX 5-1, a source familiar from Section 5.2. In both panels, the horizontal branch is represented by a dotted line, the normal branch by a solid line, and the flaring branch by a dashed line.
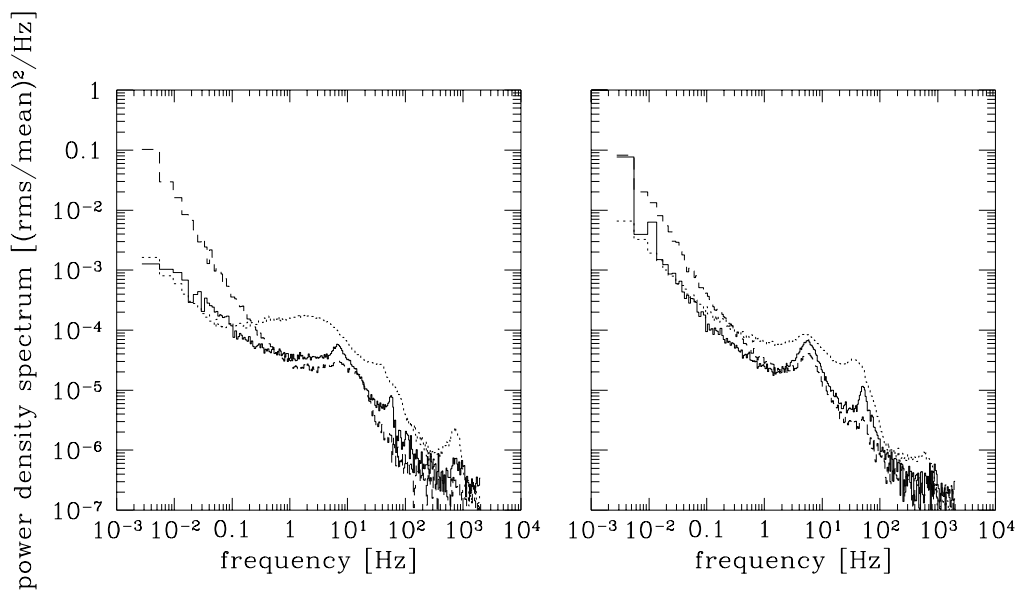

Figure 7-2: Comparison of PDSs from the three Z source branches for GX 17+2 (left panel) and GX 5-1 (right panel). The branches are represented by the following line types: horizontal (dotted), normal (solid), and flaring (dashed).

The horizontal branch for Z sources tends to exhibit more high-frequency power than the other two branches. On the other hand, the flaring branch possesses slightly more low-frequency power, i.e. for $f < 0.1$ Hz. Nevertheless, there is more qualitative similarity among Z source branches than among BHB states. In particular, QPOs

are present in similar locations for all three branches.

We now compare the PDSs of BHBs with those of NSBs. With three BHB X-ray states and three Z source branches, there are certain combinations of states and branches where NSBs consistently have greater high-frequency power than BHBs. One such combination is the BHB thermal state and the Z source horizontal branch, shown in Figure 7-3. The left panel compares XTE J1550-564 with GX 17+2, while the right panel compares GRO J1655-40 with GX 5-1. The BHBs are represented using solid lines and the NSBs using dashed lines. In both cases, the NSB PDSs are clearly higher and contain more features above 10 Hz than the BHB PDSs. Figure 7-3 supports the conclusion of Sunyaev & Revnivtsev (2000) for this combination of states.
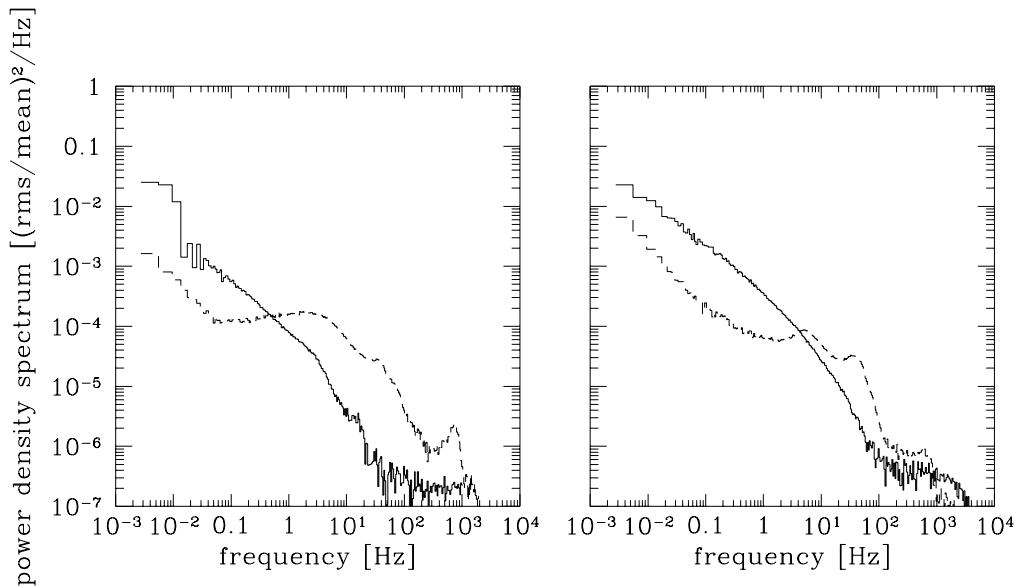

Figure 7-3: Comparison of PDSs from BHBs in the thermal state (solid) and NSBs in the horizontal branch (dashed): XTE J1550-564 vs. GX 17+2 (left panel) and GRO J1655-40 vs. GX 5-1 (right panel).

However, there also exist combinations of states where the comparison is inconclu-

sive. Figure 7-4 compares the same set of BHBs and NSBs as in Figure 7-3, but with the BHBs in the SPL state. As before, XTE J1550-564 and GX 17+2 are depicted in the left panel, GRO J1655-40 and GX 5-1 are depicted on the right, the BHBs are represented by solid lines, and the NSBs by dashed lines.
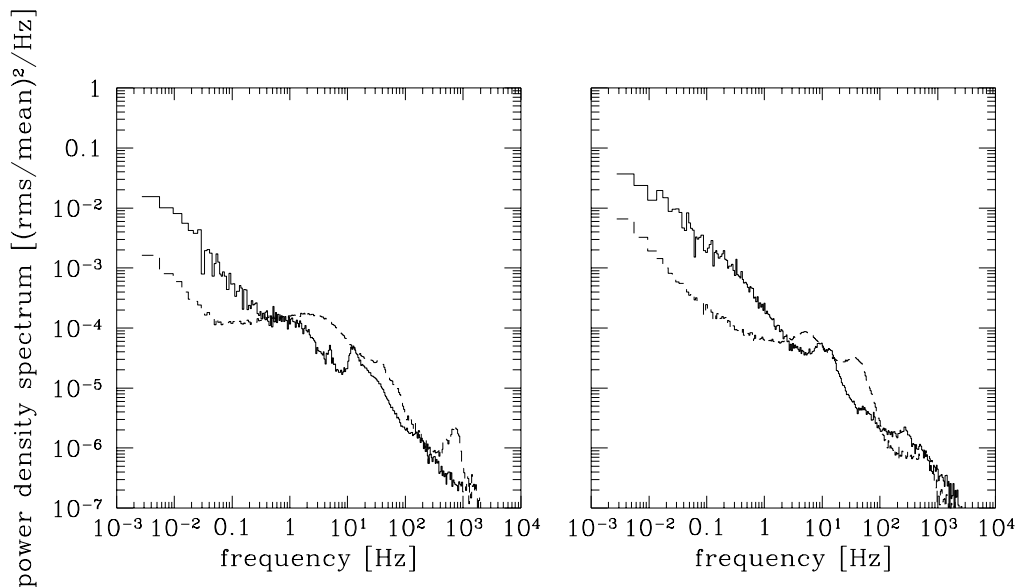

Figure 7-4: Comparison of PDSs from BHBs in the SPL state (solid) and NSBs in the horizontal branch (dashed): XTE J1550-564 vs. GX 17+2 (left panel) and GRO J1655-40 vs. GX 5-1 (right panel).

It is very difficult to distinguish between the BHB and NSB PDSs in Figure 7-4. Both contain features and significant continuum power out to at least 1 kHz and possibly beyond. We conclude therefore that the high-frequency cut-off criterion of Sunyaev & Revnivtsev (2000) cannot be extended to the SPL state of BHBs. More generally, any comparison between the PDSs of BHBs and NSBs needs to be cognizant of the X-ray state of the BHB and, to a lesser extent, the state of the NSB. Conclusions drawn from such comparisons may have to be qualified in ambiguous cases as exemplified by Figure 7-4. Since state definitions require considerations of

energy spectra and/or evolution of timing properties on a colour-intensity diagram, we conclude that BHB versus NSB classifications in general cannot be ascertained solely on the basis of PDS measurements.

Bibliography

Abramowicz, M. A., & Kluźniak, W. 2001, A&A, 374, L19

Belloni, T., Klein-Wolt, M., Méndez, M., van der Klis, M., & van Paradijs, J. 2000, A&A, 355, 271

Chakrabarti, S. K., & Manickam, S. G. 2000, ApJ, 531, L41

Dhawan, V., Mirabel, I. F., & Rodriguez, L. F. 2000, ApJ, 543, 373

Evans, R. D. 1955, Atomic Nucleus (New York: McGraw-Hill)

Fender, R. P. 2006, Jets from X-ray binaries, ed. W. H. G. Lewin & M. van der Klis, Compact Stellar X-ray Sources (Cambridge: Cambridge University Press), 381

Fender, R. P., Corbel, S., Tzioumis, T., McIntyre, V., Campbell-Wilson, D., Nowak, M., Sood, R., Hunstead, R., Harmon, A., Durouchoux, P., & Heindl, W. 1999, ApJ, 519, L165

Jahoda, K., Markwardt, C. B., Radeva, Y., Rots, A. H., Stark, M. J., Swank, J. H., Strohmayer, T. E., & Zhang, W. 2006, ApJ Supp

Jahoda, K., Stark, M. J., Strohmayer, T. E., Zhang, W., Morgan, E. H., & Fox, D. 1999, Nuclear Physics B Proceedings Supplements, 69, 210

Jernigan, J. G., Klein, R. I., & Arons, J. 2000, ApJ, 530, 875

Kubota, A., Ebisawa, K., Makishima, K., & Nakazawa, K. 2005, ApJ, 631, 1062

Kubota, A., & Makishima, K. 2004, ApJ, 601, 428

Leahy, D. A., Darbro, W., Elsner, R. F., Weisskopf, M. C., Sutherland, P. G., Kahn, S., & Grindlay, J. E. 1983, ApJ, 266, 160

Lewin, W. H., van Paradijs, J., & van der Klis, M. 1988, Space Sci. Rev., 46, 273

Makishima, K., Maejima, Y., Mitsuda, K., Bradt, H. V., Remillard, R. A., Tuohy, I. R., Hoshi, R., & Nakagawa, M. 1986, ApJ, 308, 635

McClintock, J. E., & Remillard, R. A. 2006, Black hole binaries, ed. W. H. G. Lewin & M. van der Klis, Compact Stellar X-ray Sources (Cambridge: Cambridge University Press), 157

Morgan, E. H. 2006, personal communication

Morgan, E. H., Remillard, R. A., & Greiner, J. 1997, ApJ, 482, 993

Muno, M. P., Morgan, E. H., & Remillard, R. A. 1999, ApJ, 527, 321

Nobili, L., Turolla, R., Zampieri, L., & Belloni, T. 2000, ApJ, 538, L137

Oppenheim, A. V., Schafer, R. W., & Buck, J. R. 1999, Discrete-Time Signal Processing, 2nd edn. (Upper Saddle River, NJ: Prentice-Hall)

Psaltis, D. 2006, Accreting Neutron Stars and Black Holes: A Decade of Discoveries, ed. W. H. G. Lewin & M. van der Klis, Compact Stellar X-ray Sources (Cambridge: Cambridge University Press), 1

Remillard, R. A., & McClintock, J. E. 2006, Annual Review of Astronomy and Astrophysics, 44, in press

Remillard, R. A., McClintock, J. E., Orosz, J. A., & Levine, A. M. 2006, ApJ, 637, 1002

Remillard, R. A., Muno, M. P., McClintock, J. E., & Orosz, J. A. 2002, ApJ, 580, 1030

Revnivtsev, M., Gilfanov, M., & Churazov, E. 2000, A&A, 363, 1013

Shakura, N. I., & Sunyaev, R. A. 1973, A&A, 24, 337

Shapiro, S. L., & Teukolsky, S. A. 1983, Black Holes, White Dwarfs and Neutron Stars: The Physics of Compact Objects (New York: Wiley)

Sobczak, G. J., McClintock, J. E., Remillard, R. A., Cui, W., Levine, A. M., Morgan, E. H., Orosz, J. A., & Bailyn, C. D. 2000, ApJ, 531, 537

Stirling, A. M., Spencer, R. E., de la Force, C. J., Garrett, M. A., Fender, R. P., & Ogley, R. N. 2001, Monthly Notices of the Royal Astronomical Society, 327, 1273

Sunyaev, R., & Revnivtsev, M. 2000, A&A, 358, 617

Swank, J. H. 1998, Nuclear Physics B Proceedings Supplements, 69, 12

Tagger, M., & Pellat, R. 1999, A&A, 349, 1003

Tennant, A. F. 1987, Monthly Notices of the Royal Astronomical Society, 226, 963

Titarchuk, L., & Osherovich, V. 2000, ApJ, 542, L111

van der Klis, M. 1989, Timing Neutron Stars, ed. H. Ogelman & E. P. J. van den Heuvel (Kluwer Academic Publishers), 27

—. 2006, A review of rapid X-ray variability in X-ray binaries, ed. W. H. G. Lewin & M. van der Klis, Compact Stellar X-ray Sources (Cambridge: Cambridge University Press), 39

Vignarca, F., Migliari, S., Belloni, T., Psaltis, D., & van der Klis, M. 2003, A&A, 397, 729

Vikhlinin, A., Churazov, E., & Gilfanov, M. 1994, A&A, 287, 73

Woosley, S. E., Heger, A., & Weaver, T. A. 2002, Reviews of Modern Physics, 74, 1015

Zdziarski, A. A., & Gierliński, M. 2004, Progress of Theoretical Physics Supplement, 155, 99

Zhang, W., Jahoda, K., Swank, J. H., Morgan, E. H., & Giles, A. B. 1995, ApJ, 449, 930

Zhang, W., Morgan, E. H., Jahoda, K., Swank, J. H., Strohmayer, T. E., Jernigan, J. G., & Klein, R. I. 1996, ApJ, 469, L29