THE ASTROPHYSICAL JOURNAL, 400:138–152, 1992 November 20

THE X-RAY VARIABIITY OF NGC 6814: POWER SPECTRUM

C. DONE,[1,2] G. M. MADEJSKI,[3] R. F. MUSHOTZKY, AND T. J. TURNER[3]
NASA/Goddard Space Flight Center, LHEA Code 660, Greenbelt, MD 20771
AND
K. KOYAMA AND H. KUNIEDA
Department of Astrophysics, Nagoya University, Chikusa-Ku, Nagoya 464, Japan
Received 1992 February 18; accepted 1992 May 28

ABSTRACT

NGC 6814 is unique among Seyfert galaxies in showing strong evidence for periodic X-ray behavior, first detected from an *EXOSAT* "long-look observation." Here simulation techniques are used to obtain the X-ray variability power spectrum of the unevenly sampled *Ginga* light curve of the object. A simple power law is not an adequate description of the power spectrum, with the residuals showing excess power on time scales consistent with the periodicity of 12,100 ± 100 s seen in the *EXOSAT* data. Over 75% of the total source variability is due to the periodic component (rms amplitude of 36%). The residual variability can be described as the more usual "flicker noise" $f^{-1.1}$ power law. This continues unbroken to time scales shorter than 300 s, so that the region from which the variability arises must be smaller than 10^{13} cm.

Folding techniques, combined with the simulations to obtain the error, give the period as 12,132 ± 3 s. While the period is consistent with that seen in the *EXOSAT* data, the shape of the folded light curve is very different. There are three main peaks, two of which are separated by an extremely sharp dip, instead of the single peak and small harmonic structure observed by *EXOSAT*. Using the dip as a fiducial mark, a second *Ginga* observation of the source taken 1 year later is found to be consistent with being completely periodic and phase-coherent with the first *Ginga* observation. Thus the period is consistent with being constant over a time span of 6 years, but phase coherence is only maintained on time scales of ≤ 1 year.

Subtle differences in the shape of the folded light curve with energy, and the very large amount of power in the periodic component, suggest occultation as its origin, although amplification of variability from an X-ray–emitting "hot spot" through gravitational lensing is also possible. The former suffers from the very arbitrary nature of the periodic time scale, while the latter is unattractive because it is difficult to trivially produce the different shape of the folded light curve seen in the *Ginga* and *EXOSAT* data. That both these models probably fail to provide an adequate explanation may be due to the added complexity of anisotropic X-ray emission, suggested by the discrepancy between the lack of soft photons implied by the flat spectrum and the copious source of soft photons available from reprocessing in the iron-line–producing material.

Subject headings: galaxies: active — galaxies: individual (NGC 6814) — methods: data analysis — X-rays: galaxies

1. INTRODUCTION

X-ray emission from active galactic nuclei (AGNs) provides a large contribution (~10%) to the total bolometric luminosity of these objects, suggesting that the X-rays are closely associated with the primary energy generation source. This emission is also highly variable on short time scales, which by light crossing time arguments implies a very small size scale for the emission region. The efficiency of the X-ray emission provides the strongest evidence in support of the standard theory that all the activity seen is ultimately powered by accretion onto a black hole. No other spectral region in radio-quiet AGNs varies so rapidly, indicating that the X-rays are a probe of the smallest size scales and hence of the environment and energy generation processes close to the central engine.

EXOSAT provided the best data with which to study X-ray variability in AGNs, since the satellite was able to observe a source continuously for up to 3 days. Power spectral analysis could then be used to determine the statistical nature of the variability. This showed that power was present on all time scales, from 10^3 s (where the source signal was typically lost in the Poisson noise) to days. A power-law fit to the power spectrum, $P(f)$, was adequate in most cases, with $P(f) \propto f^{-\alpha}$, $\alpha \sim 1$ (Lawrence et al. 1987; McHardy & Czerny 1987; McHardy 1990). In such power spectra there is no characteristic time scale, showing that over the observed frequency range the variability is self-similar and scale-invariant, and that the fastest variability seen is dependent on the length of the observation and sensitivity of the detector. A faster event may be seen by simply observing over a longer period, so size scales deduced from "two-folding time scales" are not a characteristic of the source (Press 1978), although they still imply an upper limit to the size scale.

A power spectrum of strictly f^{-1} form cannot continue indefinitely since the total power is unbounded at high and low frequencies. A high-frequency steepening is expected, which would provide the elusive characteristic frequency, which can then be translated to a unique size scale, giving the dimension of the smallest regions contributing to the X-ray emission. There are indications of such a high-frequency rollover in the

[1] NAS/NRC Resident Research Associate.
[2] Present address: Department of Physics and Astronomy, University of Leicester, University Road, Leicester LE1 7RH, UK.
[3] Universities Space Research Association.

EXOSAT data (McHardy 1990), although the statistical significance of this is hard to assess. Data with better signal-to-noise are required to determine the high-frequency behavior of the power spectrum unambiguously. In NGC 4051, the one object where this has been attempted, data from the more sensitive *Ginga* satellite show that there is little evidence for high-frequency steepening on time scales above 100 s, where again the noise dominates (McHardy 1990).

There are a few exceptional objects with power spectra that are not well fitted by a power law. *EXOSAT* data from NGC 6814 (Mittaz & Branduardi-Raymont 1989; Fiore, Massaro, & Barone 1992b) and NGC 4151 (Fiore et al. 1989) have been claimed to show a periodic component. While the evidence for this is fairly weak in the case of NGC 4151, in NGC 6814 there is a very significant feature in the power spectrum for the longest observation associated with a periodic flare in the light curve (Mittaz & Branduardi-Raymont 1989; Fiore et al. 1992b), and features associated with harmonics from this period are present in all of the other four shorter observations (Fiore et al. 1992b). The reality of the periodic component in NGC 6814 then seems convincing, although its nature is not understood.

NGC 6814 is also unusual in other ways. It is one of the most rapidly variable AGNs, with large-amplitude changes on times scales as short as 100 s seen by Tennant & Mushotzky (1983) in the *HEAO 1* A-2 instrument. This extreme variability was not observed in the five separate observations of NGC 6814 with *EXOSAT*, although they did show the periodic component to the variability as described above, and strong changes in the X-ray–absorbing column (Mittaz & Branduardi-Raymont 1989).

A 3 day *Ginga* observation of NGC 6814 showed the recurrence of the extreme variability seen in *HEAO 1*. Kunieda et al. (1990) show that the source decreases by a factor of 2 in luminosity over 50 s. The spectrum shows many peculiarities in terms of the "canonical" spectra of AGNs (Pounds et al. 1990; Zdziarski et al. 1990). The spectral index is very flat, with $\alpha = 0.4$; the iron-line equivalent width is very large (~ 300 eV for a narrow Gaussian: Kunieda et al. 1990; Turner et al. 1992); and the spectrum shows little evidence for a pronounced high-energy excess associated with reflection of X-rays from the accreting material (Turner et al. 1991; Turner et al. 1992). A comparison of the line and continuum emission, made possible by the very rapid variability and the strong iron line, showed no detectable lag on time scales longer than 256 s (the shortest time over which spectral fitting to derive the line flux was possible; Kunieda et al. 1990). This implies that the difference in light travel path between the reprocessed flux and the continuum emission is less than 256 s.

The intense and rapidly variable iron line points to there being substantial amounts of material close to the X-ray source. However, little sign of this material can be seen in terms of a Compton reflection spectrum expected from material out of the line of sight (Lightman & White 1988; Guilbert & Rees 1988) or in terms of a simple edge expected if the material had a more spherical distribution. However, Fe Kβ and Ni Kα emission associated with the intense Fe Kα line can mask some of the edge expected from reflection (Turner et al. 1992), possibly allowing a consistent solution to be found where the line is produced in a highly ionized accretion disk.

Here the variability of these *Ginga* data is examined in depth. The analysis is complicated by the uneven sampling inherent in observations taken from low Earth orbit over a period of days. In a Fourier analysis, this has the effect of convolving the Fourier transform of the data with that of the sampling pattern. Numerical simulations are used to constrain the intrinsic power spectral shape, through a comparison of the observed (convolved) power spectrum of the data with that produced from a simulated light curve, of known power spectral form, sampled in the same uneven way as the *Ginga* data.

The data are described in § 2, and the problems of uneven sampling are briefly discussed in § 3. Simulations of light curves with simple power-law power spectra are shown in § 4 to be a poor fit to the observed power spectrum, with the residuals showing peaks in the same frequency range as the periodicity seen in the *EXOSAT* data. A folded light curve analysis shows that the periodic component is indeed present in this new data set. While the shape of the folded light curve shows very different structure from that seen in the *EXOSAT* observations, the period is identical within the errors quoted. Thus it seems that the period is stable over time scales of 6 years. Phase coherence is not maintained over all this time although the data are consistent with there being phase coherence between this observation and a second *Ginga* observation taken a year later. Apart from the periodic component, the power spectrum is well described by an $f^{-1.1}$ shape, with no evidence for a high-frequency break on time scales longer than 300 s. The results are discussed in § 5, where current theories for the origin of the periodicity are examined critically, and a summary of the paper is given in § 6.

2. THE DATA

Details of the *Ginga* satellite configuration are given in Turner et al. (1989). The data selection and background subtraction procedures used only data from the top-layer *Ginga* detectors, collected in the low-background "remote" orbits (see Appendix A for details). This results in a light curve of 1848 points, intrinsically binned on 16 s time intervals, which is then divided into four energy bands, namely, very low (VL; 1.67–2.80 keV), low (L; 2.80–5.69 keV), medium (M; 5.69–6.84 keV), and high (H; 6.84–20.90 keV) energy. Details of source count rates in each band are given in Table 1. These bands were chosen as being most likely to represent different components. Changes in the hardness ratio between the VL and L energy bands (see Turner et al. 1992) are consistent with those expected if a warm (partially ionized) absorber is also present in the system, affecting the variability in the VL band. Line emission from Fe Kα is known from the spectrum (Kunieda et al. 1990; Turner et al. 1992) to contribute mostly to the M band, and all higher energy data are added together to get a good signal-to-noise ratio in the H continuum.

TABLE 1

COUNT-RATE STATISTICS

Energy Band	$\bar{x}$	$\bar{x}_{\rm weighted}$	$\sigma^2_{\rm obs}$	$\sigma^2_{\rm err}$	$\sigma^2_{\rm int}$	$\sigma_{\rm int}/\bar{x}_{\rm weighted}$
VL	0.850	0.727	0.832	0.413	0.418	0.889
L	5.87	5.52	7.26	1.26	6.00	0.443
M	2.20	2.01	1.20	0.400	0.800	0.445
H	4.47	4.49	4.46	1.17	3.29	0.404
VH.....	0.167	0.0865	0.361	0.309	0.0527	2.65

NOTE.—Shown are the average, error-weighted average, observed variance, error bar variance, intrinsic variance, and intrinsic rms variability for each energy band.

3. UNEVENLY SAMPLED DATA AND FOURIER ANALYSIS

Power spectra are the most intuitively obvious of the Fourier techniques, showing the relative amount of variability at each frequency considered. However, Fourier analysis of the X-ray variability of NGC 6814 in observations with the *Ginga* satellite is complicated by the uneven sampling of the light curve. On time scales longer than the orbital period of the satellite (96 minutes) the observation is interrupted by Earth occultation of the source. There are also periods of high instrument background as the orbit precesses through the South Atlantic Anomaly. This gives a characteristic sampling pattern shown in Figure 1, where the zero point is 1989 April 4 02:50:02.943 (UT). Standard fast Fourier transform techniques to compute power spectra and auto- and cross-correlations cannot be used without interpolating between the unevenly spaced data and resampling on an evenly spaced grid. This is unsatisfactory in general, since it assumes the source behavior over periods when it is unobserved, and in particular because there are more gaps than data in the *Ginga* sampling pattern. Algorithms which make use only of the measured data points have been developed for power spectral analysis (Deeming 1975; Scargle 1982; Lomb 1976), but while these avoid the problems of interpolation, the resultant power spectra are distorted because of the uneven sampling and so cannot be simply interpreted.

A Fourier transform of sampled data is a convolution for the Fourier transform of the intrinsic light curve with that of the sampling pattern (Deeming 1975). Where the variability is intrinsically stochastic, the observed power spectrum is given by the convolution of the intrinsic power spectrum with the square of the transform of the sampling pattern (Deeming 1975), henceforth called the window function. This convolution can be ignored if the window is a δ-function. This is never the case with a finite data train, but with discrete, evenly sampled data the window function has the same effect as a δ-function if the power spectrum is calculated only at the discrete set of frequencies $\{1/T, 2/T, \ldots, 2/(\Delta t)\}$, where T is the length of the observation and Δt is the binning time (see, e.g., Deeming 1975). Power calculated at these frequencies is then independent of the calculation of the power at all other frequencies, so an error bar can be obtained by a rms error from rebinning (usually logarithmically) the power spectrum.

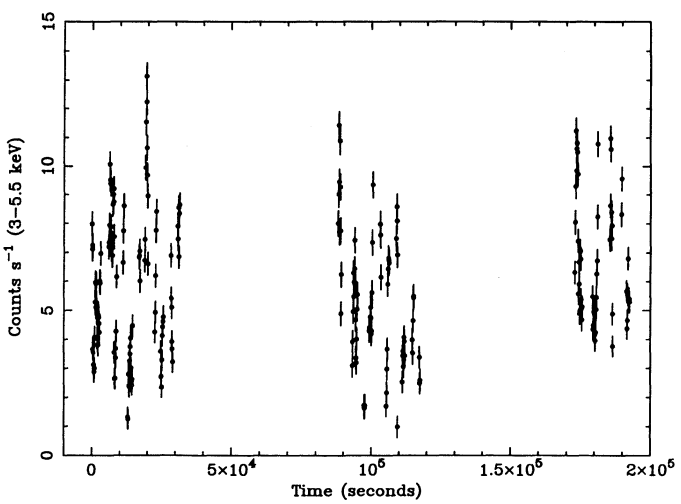

FIG. 1.—L light curve of NGC 6814, binned on 128 s, showing the sampling pattern for the observation. Earth occultation of the source gives rise to the short gaps, and the two large gaps are due to orbit precession through the South Atlantic Anomaly. Zero time is 1989 April 4 02:50:02.943 UT.

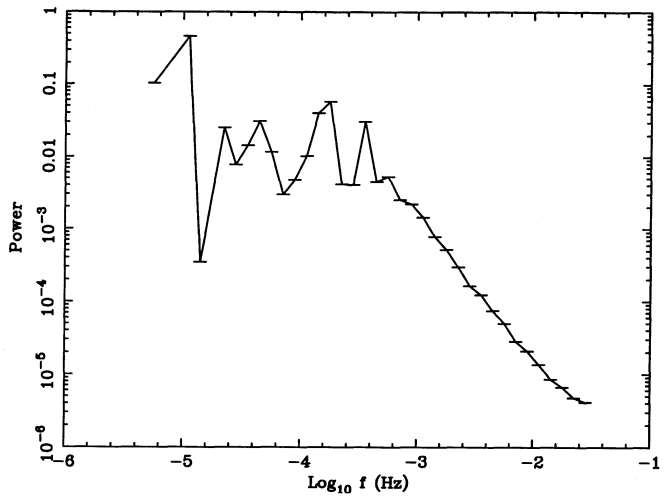

FIG. 2.—Squared modulus of the Fourier transform of the sampling pattern of the data shown in Fig. 1. This is the window function (calculated using a slow Fourier transform rather than the Lomb and Scargle algorithm), through which the source power spectrum is convolved. Note the strong peak at times corresponding to 96 minutes (the orbital time scale) and its first harmonic, and the flattening to low frequencies.

For unevenly sampled data the effects of the convolution cannot be avoided. This distorts the shape of the power spectrum with leakage of power over a wide frequency band. Hence the powers calculated on the frequency set $\{1/T, 2/T, \ldots, 2/(\Delta t)\}$ are no longer independent, and an rms error is an underestimate of the true uncertainty on the power at each frequency (see Fiore et al. 1992a). Figure 2 shows the window function for the sampling pattern of the NGC 6814 data, calculated as for evenly sampled data, i.e., at $\{f = 1/T, 2/T, \ldots, 1/(2\Delta t)\}$, and then logarithmically rebinned (10 points per decade). Peaks at the orbital frequency (1.74×10^{-4} Hz) and its harmonics can clearly be seen, as also can the fact that this window function does not go to zero.

Where the variability is stochastic in nature, the intrinsic power spectrum of unevenly sampled data can be constrained by convolving a model form of the power spectrum through the window function. This is then compared via the χ^2 statistic with the observed rebinned power spectrum, where the error bars are corrected for the nonindependence of adjacent points in the unbinned power spectrum. Further details of this technique are given in Fiore et al. (1992a). However, for data in which a deterministic (periodic) component is suspected, this method does not apply (see Deeming 1975). Instead, numerical simulations of a variable light curve of known intrinsic power spectral form are used, sampled in the same way as the observed light curve to introduce the same distortion as in the original data. The power spectrum of the simulated light curve is then compared with that of the data to see whether the model description of the power spectrum is acceptable, where the error bars come from repeated simulations.

4. THE POWER SPECTRUM

The power spectrum of the L-band light curve was calculated using the Lomb and Scargle algorithms, as implemented by Press & Teukolsky (1988), and logarithmically rebinned to

have 10 points per decade. There are three main problems with the calculation of the power spectrum. The noise level conceals the true shape of the power spectrum at high frequencies; the calculated points in the power spectrum are not independent, so an rms error is an underestimate of the true variance (see § 3); and the whole spectrum is distorted by its convolution with the window function (see § 3).

4.1. *Noise Power*

With the first of the problems, the asymptotic value of the measurement noise power can be easily estimated from the error bars, $dx(i)$, on the N light-curve data points, $x(i)$, as

$$P_{\rm EB,noise} = \frac{(1/N)\Sigma\, dx(i)^2}{\Sigma[x(i) - \bar{x}]^2/(N-1)} = \frac{\sigma_{\rm err}^2}{\sigma_{\rm tot}^2}, \quad (1)$$

where $\bar{x}$ is the mean count rate, $\sigma_{\rm tot}^2$ is the total variance of the data, and $\sigma_{\rm err}^2$ is the variance introduced by the measurement error. Hence the intrinsic variance of the souce is $\sigma_{\rm int}^2 = \sigma_{\rm obs}^2 - \sigma_{\rm err}^2$.

The error on the noise power is also an important quantity, since it significantly increases the error bars on the power at high frequency. The $100(1 - a)\%$ confidence range in the true variance, $\sigma_{\rm true}^2$, for an observed variance of σ is (see e.g., Hoel 1984)

$$\frac{(N-1)\sigma^2}{\chi^2_{N-1,1-a/2}} \le \sigma_{\rm true}^2 \le \frac{(N-1)\sigma^2}{\chi^2_{N-1,a/2}}. \quad (2)$$

For $N = 1848$, a 1 σ error bar (i.e., 67% confidence, $a = 0.33$) gives $0.968\sigma^2 \le \sigma_{\rm true}^2 \le \sigma_{\rm true}^2 \le 1.033\sigma^2$. This fractional error, applied to both $\sigma_{\rm err}^2$ and $\sigma_{\rm tot}^2$, gives a 1 σ uncertainty in the noise power such that $P_{\rm EB,noise} = 0.170 \pm 0.028$.

To check that there is no significant systematic error adding to this measurement noise, the noise is estimated in a second, independent way. The power spectrum at high frequencies where the measurement error dominates over the signal can be approximated by a constant, giving $P_{\rm HF,noise} = 0.185 \pm 0.005$. While this uses the error bars obtained from rebinning, the hence the error on the noise is an underestimate, the consistency between the two estimates of the measurement noise power suggests that the error bars assigned to the light curve are an accurate representation of the measurement error and that there is no significant systematic error. This point is examined in more detail in Appendix A. Since the first method gives the most conservative error, $P_{\rm EB,noise}$ is used in all the following calculations.

4.2. *Power Spectrum Simulations*

Much more serious than the noise are the problems produced by the underestimate of the error bar and the distortion of the data through the window function. Here simulations of light curves of known power spectrum, with the same mean, variance, measurement error, and sampling characteristics (which includes the effects of the *Ginga* window and the intrinsic binning of the light curve into 16 s intervals) as the real *Ginga* data, are used as a means of investigating the form of the intrinsic power spectrum. Appendix B describes in detail how the simulated light curves are produced.

For each trial form of the intrinsic power spectrum, 10 (the minimum number of points in order for the assumption of the normal distribution to be justified) random simulations of light curves with that form of power spectrum, scaled to the NGC 6814 data and with the same uneven *Ginga* sampling window,

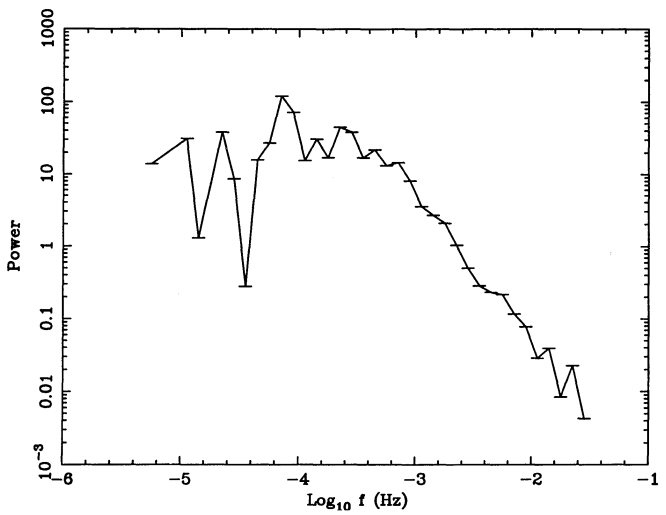

FIG. 3.—Power spectrum of the *L*-band light curve, calculated with the Lomb and Scargle algorithm using only the measured data points. While this is free from errors arising from interpolation, it is not free from the convolution with the window function (note the flattening at low frequency). The power spectrum is rebinned logarithmically and Poisson noise subtracted.

are produced. The calculated power spectrum of each simulated light curve is logarithmically rebinned to 10 points per decade. A total power spectrum from the simulation is then found by stacking these rebinned power spectra together, so that each point is the average of 10 independent points. The error on the power at each frequency is then well defined as the rms error on the mean of the 10 independent estimates of the power. The measurement noise power (§ 4.1) was then subtracted from this average power spectrum and the error bars increased accordingly.

The *Ginga* data from NGC 6814 were then used as a particular realization of the hypothesis that the power spectrum could be characterized by the assumed $P(f)$ shape. Its rms error bars were discarded as being misleading (see § 3), giving the power spectrum realization shown in Figure 3 after noise subtraction. The goodness of fit is calculated from the well-determined error bars on the total simulated power spectrum. While this approach seems unusual, the definition of the χ^2 parameter uses the variance of the model population from which the data are drawn, and the error bars on the data are used to approximate this. So using the error bars from the model is actually the more correct way to calculate χ^2 and hence assess the probability that the hypothesis of the underlying shape of the power spectrum is correct.

These simulations were calculated for an intrinsic power spectrum $P(f) \propto f^{-\alpha}$, $\alpha = \{1.0, 1.2, 1.4, 1.6, 1.8\}$. With this red-noise power spectrum, the total variance is determined by the power at the lowest frequencies. Figure 4*a* shows a plot of χ^2 versus the power-law slope for each of these. Each power spectrum has 35 frequency bins, but the model uses the average, variance, and error bar variance as calculated from the real data, so the number of degrees of freedom is reduced to 32. For the model to be acceptable at even the 1% level requires $\chi^2 \le 53.5$, whereas the best-fit model has $\chi^2 = 183$ for a slope of $\alpha = 1.4$.

Since the total variance of the source is also subject to uncertainty in the same manner as the noise variance, the normalization of the model power spectrum was varied to see whether

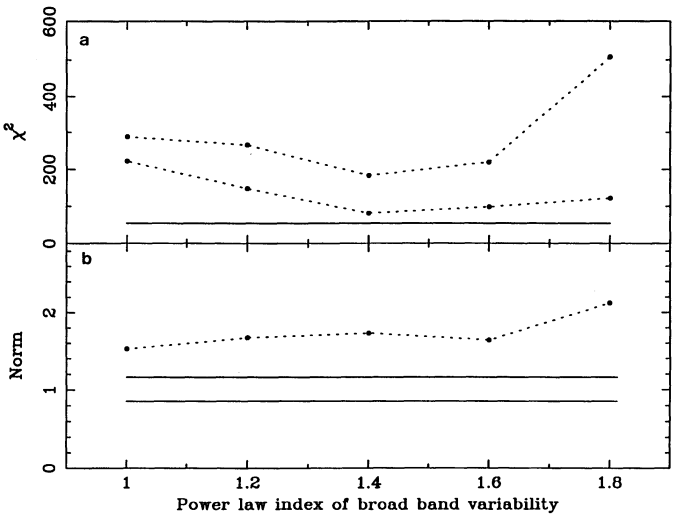

FIG. 4.—The upper curve in (*a*) shows the goodness-of-fit parameter, χ^2, plotted against the trial power-law power spectral index for a simulated light curve with the same mean variance, error bar properties, and sampling as the real L-band data from NGC 6814. The lower curve shows the minimum χ^2 obtained through allowing the variance of the simulated light curve (i.e., the normalization of the power spectrum) to vary. None of these are acceptable at even the 1% confidence level (which would require $\chi^2 \leq 53.5$). The predicted 3 σ variance of the mean corresponds to a relative normalization change of 1.0 ± 0.15. This allowed region is shown by the solid line in (*b*). None of the normalization changes are consistent with this, showing that a simple power-law power spectrum is not an acceptable description of the variability from NGC 6814.

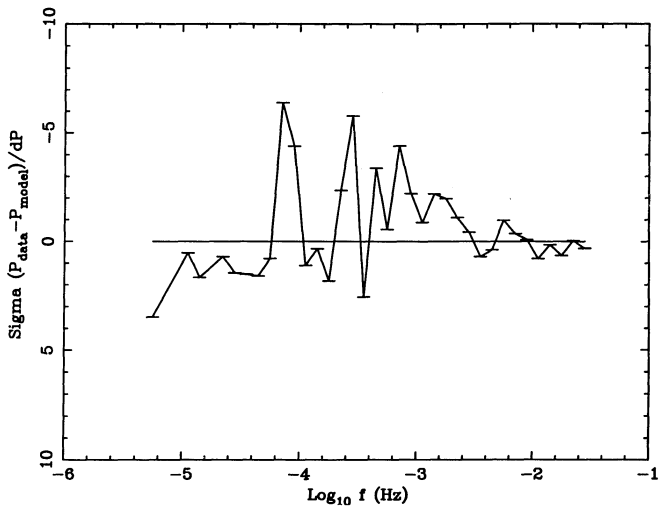

FIG. 5.—Residuals from the best-fit power-law power spectrum ($\alpha = 1.4$) with normalization fixed. These do not show a systematic offset but rather show a series of sharp peaks. The lowest frequency peak is the most significant and corresponds to time scales of 10^4 to 1.5×10^4 s, overlapping the time scale of 12,200 $\pm$ 100 s at which the *EXOSAT* data showed evidence for a periodic phenomenon.

χ^2 could be minimized further. Figure 4*a* shows the new minimum found as the lower line, and the normalization used to obtain that minimum is shown in Figure 4*b*. The 1 σ error on the intrinsic variance is $\sqrt{2} \times 0.035$, i.e., only 5%, but the change in normalization required for the minimum is always greater than 50%.

A closer look at the residuals from the fits clarifies the situation. The residuals of the data from the $\alpha = 1.4$ power-law model with normalization fixed to the variance of the original data set are shown in Figure 5, and it can be seen that the main contributions to χ^2 are from sharp features rather than from a systematic discrepancy in the continuum power. They peak in a frequency bin spanning 10^{-4} to $10^{-4.2}$ Hz, i.e., including the time scale of ~12,100 s claimed as a period by Mittaz & Branduardi-Raymont (1989) from *EXOSAT* data. The excess power in the *Ginga* data cannot be due to the *Ginga* sampling, since the simulations also include the same sampling pattern.

An excess of power over that predicted by the model power law over a narrow range of frequencies in the power spectrum distorts the distribution of the residuals. A better fit can then be found by increasing the variance of the model data, as this increases the observed power at all frequencies, making the peak less significant. This then explains the unphysically large change in normalization required to find the minima, which, together with the fact that the χ^2 is never acceptable even at the 1% level, shows that the data are not well described by a variability power spectrum which is a featureless power law. Clearly, then, there is some other contribution to the power spectrum which is probably deterministic (periodic) rather than random. The next sections attempt to characterize this additional component more precisely.

4.3. *The Folded Light Curve*

Fourier techniques are most sensitive to a periodic component of sinusoidal form, but the periodicity detected in the *EXOSAT* data is definitely nonsinusoidal in shape (Mittaz & Branduardi-Raymont 1989; Fiore et al. 1992b). Folding techniques are then more appropriate since they are less affected by harmonics. In the simplest form of these techniques, the light curve is folded on the test period, and the χ^2 statistic is used to determine whether the folded light curve is consistent with a constant value. The worse the fit to a constant, the more probable that the light curve does contain a periodic component on that time scale.

This technique, described by Leahy et al. (1983) and extended by Davies (1990, 1991) is used to search the period range of $P_{min} = 10{,}000$ s to $P_{max} = 15{,}840$ s (implied by the excess power between 10^{-4} and $10^{-4.2}$ Hz) in steps of 1 s. The light curve, folded on the trial period P, is binned into M pulse phase bins. The null hypothesis (no periodic component) is that the folded light curve is uniform with phase, i.e., is consistent with a constant. Leahy et al. (1983) define a goodness-of-fit statistic, Q^2, as

$$Q^2 = \sum_{j=1}^{M} \frac{[\overline{xf(j)} - \bar{x}]^2}{\sigma^2_{xf}(j)}, \quad (3)$$

where $\overline{xf(j)}$ is the average count rate in phase bin j, with variance $\sigma^2_{xf}(j)$. If there is no significant deviation from a constant, then $\sigma^2_{xf}(j) \approx \sigma^2_{tot}/n(j)$, so that

$$Q^2 = \frac{1}{\sigma^2_{tot}} \sum_{j=1}^{M} n(j)[\overline{xf(j)} - \bar{x}]^2, \quad (4)$$

and this Q^2 statistic is approximately χ^2 distributed with $M - 1$ degrees of freedom. However, for this approximation to be valid, it is required that there be at least 10 points in each phase bin, i.e., that $N \gg M$ (Davies 1990). This is not generally the case, especially with unevenly sampled data where there can be a large difference in the number of points in each phase bin.

A better test statistic is given by the L-statistic (Davies 1990), defined as

$$L = \frac{(N - M)Q^2}{(M - 1)(N - 1 - Q^2)}. \quad (5)$$

In contrast to the Q^2 statistic, this has an F-distribution, with $M - 1$ and $N - M$ degrees of freedom, for all N and M. It is an exact rather than an approximate statistic and so is more sensitive (Davies 1990).

4.3.1. *The Significance of the Periodic Component*

As stated by Davies (1991), the period with the largest value of the L-statistic is the one least likely to be reconciled with the null hypothesis and so is the best candidate for the true period. Figure 6 shows the variation L with trial period, with a clear peak at $\sim$12,100 s. Fitting the peak in L with a Gaussian over a restricted range (12,110–12,150 s) of fold periods gave $P = 12{,}130.9$ s, with $L_{\rm peak} = 40.25$. A formal error on the centroid of the Gaussian gives an uncertainty on the period of ± 0.6 s, but this is an underestimate, first because the points fitted with the Gaussian are not independent, and second because the fitting procedure does not correct for any possible distortion in L introduced by the effects of uneven sampling of red-noise power. This period is consistent with that found from the *EXOSAT* data, to within their quoted errors of 12,200 $\pm$ 100 s (Fiore et al. 1992b).

The probability of obtaining such a large value of L is extremely small. Using the normal distribution as an approximation to the F-distribution (Abramowitz & Stegun 1972, § 26.6.14) this formally corresponds to a 41.5 σ result, or a probability of less than 10^{-367}. Assessing the *true* significance of this peak is more difficult. First, the maximum in L does not represent the standard case where a period is superposed on uncorrelated (white) noise. The underlying broad-band power seriously distorts the probabilities. This can be seen in Figure 6 because the value of L is typically much larger than unity, i.e., naively, there is a significant period at most of the fold times examined. This is analogous to the problem of assessing the probability of a peak in the power spectrum in the presence of red noise, where the solution is to calculate the probability on the ratio of the signal-to-noise power rather than on the signal probability alone (Scargle 1982). Schwarzenberg-Czerny (1989) suggest that the false-alarm probability can be estimated by dividing the peak L by the mean L over the period range studied. This then gives the probability of seeing the observed maximum in L in the presence of red noise as 1.11×10^{-33} ($L = 4.10$).

Second, the test examined $N_p = 5840$ possible periods. While not all of these are independent (since the separation in frequency is much less than $1/T$ and the data are not evenly sampled), the most conservative assumption (Davies 1991) is that the probability of obtaining the peak value of L is increased by a factor N_p, i.e., is 6.5×10^{-30}.

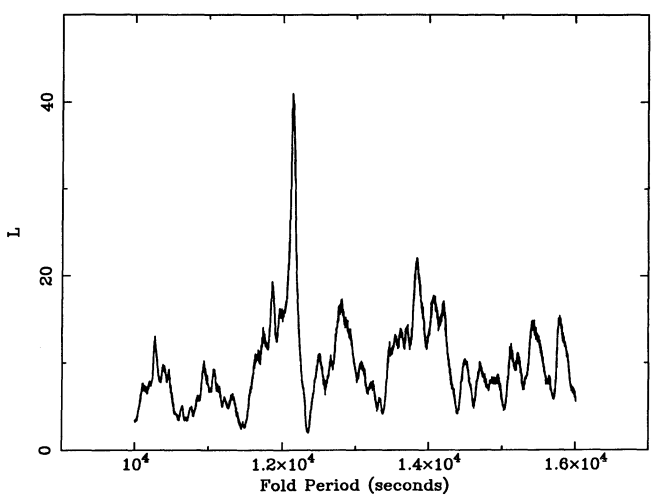

FIG. 6.—L-statistic for the set of trial periods determined from the peak of the residuals in the power spectrum as shown in Fig. 5, with a clear peak at $\sim$12,100 s.

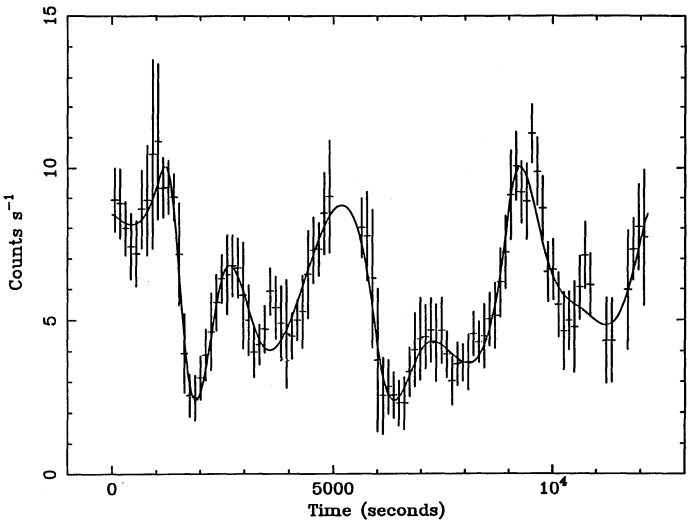

FIG. 7.—L-band light curve folded on the best-fit period of 12,130.9 s determined from the L-statistic, together with a spline fit to the variability.

4.3.2. *The Periodic Component*

Figure 7 shows the light curve folded on the best estimate of the period into 100 phase bins. The error bars are calculated from the error on the mean, with a correction for the fact that some of the points within a single bin have mutual correlations and thus are not independent (Edelson & Krolik 1988). The autocorrelation length in NGC 6814 is $\geq$100 s (Tennant & Mushotzky 1983), hence in a single phase bin of 121.309 s most consecutive points on the light curve are correlated rather than independent. Each plotted phase bin contains 10–32 points in total, of which 2–4 are independent. Overlaid on the folded light curve is a 15 point spline fit to the shape, detailed in Table 2. The zero point is 01:15:50.0 UT 1989 April 4 (the slight shift from the light curve shown in Fig. 1 is due to not normalizing the start to zero.) There are three main peaks, the shapes of which differ. In particular, $\sim$2000 s into the folded light curve, there is a very sharp dip and rise feature. This can be used as a fiducial marker, though the second dip at $\sim$6000 s is similar in having a very sharp drop (but with a somewhat longer recovery time). This contrasts sharply with the folded light curve derived from the *EXOSAT* observations, which showed only a single flare event, with some harmonic structure (Mittaz & Branduardi-Raymont 1989; Fiore et al. 1992b).

The spline fit was expanded to produce a light curve spanning the entire data set, and sampled on the same sampling pattern as the original data. This has fractional rms power ($\sigma_{\rm int}/\bar{x}$; Lewin, van Paradijs, & van der Klis 1988) of 0.36,

TABLE 2

SPLINE-FIT COEFFICIENTS

x	y	d^2y/dx^2
0000.000	8.470620	$-2.6325250 \times 10^{-5}$
1000.000	9.324980	2.2754894×10^{-5}
1300.000	9.728940	$-9.9614750 \times 10^{-5}$
1700.000	3.457260	7.6199438×10^{-5}
2500.000	6.383830	$-3.3760316 \times 10^{-5}$
3313.840	4.411830	1.4155341×10^{-5}
5861.370	5.637320	$-1.9729852 \times 10^{-5}$
6045.490	3.719710	3.4502944×10^{-5}
6999.250	4.023940	$-1.0992186 \times 10^{-5}$
8943.170	8.009970	2.1186743×10^{-5}
9000.000	8.697450	$-5.5097677 \times 10^{-5}$
9800.000	7.312060	1.2801906×10^{-5}
10700.00	5.304640	$-2.7119063 \times 10^{-6}$
11529.10	5.245580	1.3040651×10^{-5}
12130.90	8.470260	$-4.1375915 \times 10^{-6}$

which is close to the fractional rms power of the total signal (corrected for the variance introduced through measurement error) of 0.44. A ratio of variances shows that the periodic component contributes 75% of the total signal variance. This is actually an underestimate of the power in the periodic component, because of the inherent smoothing associated with a spline fit.

Figure 8*a* shows the power spectrum of the periodic component, whereas Figure 8*b* shows the power spectrum after uneven sampling in the same way as the real data. The latter shows a remarkable resemblance to the power spectrum of the real data (Fig. 3), as expected if the periodic component accounts for the majority of the power, with a poor match only at high frequencies. The strong secondary peak in the residuals from a power-law power spectrum (Fig. 4) at $\sim 2.5 \times 10^{-4}$, corresponding to ~ 4000 s, can now be seen to be due to both power at ~ 4000 s in the perodic component (see Fig. 8*a*) from the approximately equispaced triple peaks in the folded light curve, and to sidebands of the main peak from the orbital period.

4.3.3. *Energy Dependence of the Periodic Component*

A repeat of this analysis on the remaining three energy bands shows that the periodic component is clearly seen in each, with estimates of the false-alarm probability at periods of $12{,}133.0 \pm 0.8$, $12{,}132.8 \pm 1.0$, and $12{,}132.5 \pm 0.8$ s of 8.9×10^{-14}, 8.9×10^{17}, and 1.4×10^{-21} for the VL, M, and H bands, respectively (see Figs. 9*a*, 9*b*, and 9*c*). Taken together, the four periods are consistent with a constant period ($\chi^2 = 5.95$ for 3 degrees of freedom) even with these nominal error bars, which are an underestimate of the true error. An average and rms error of these four periods gives a time scale of $12{,}132.3 \pm 1.0$ s, where this error is a true estimate of the scatter about the mean (although with a small number of points) but does not take account of any systematic distortion in the time scale of the peak in L that may be present as a result of the nonuniform underlying distribution in L.

Figures 10*a*–10*c* show the folded light curve, together with the spline fit to the L band given in Table 1 (scaled by the error-weighted average count rate in the relevant band compared with that in the L band). It can be seen that the scaled spline is a good fit to the folded light curve in all bands, but that there is a hint of some energy-dependent behavior. The first sharp dip and rise is narrower in the M and H bands than predicted by the L-band spline, although it seems equal in amplitude. However, the large error bars on the folded light curves mean that statistically a hardness ratio of two bands is never required to differ from a constant. This behavior contrasts with that in the VL band, in which the dip is wider and deeper, and the folded light curve generally varies by a larger amplitude than that predicted by the scaled spline. Since the periodic component is dominated in the variability, this can be seen in the rms amplitude of variability of the raw light curves (see Table 1). While the L, M, and H bands are all consistent with the same amount of variability, the amplitude of variability in the VL band is almost twice that in the other bands.

A decrease in the width of the dip at higher energies is also seen in a more detailed analysis of the spectral variability during the dips by Leighly, Kunnieda, & Tsuruta (1992), who suggest that this is evidence for the dips being due to

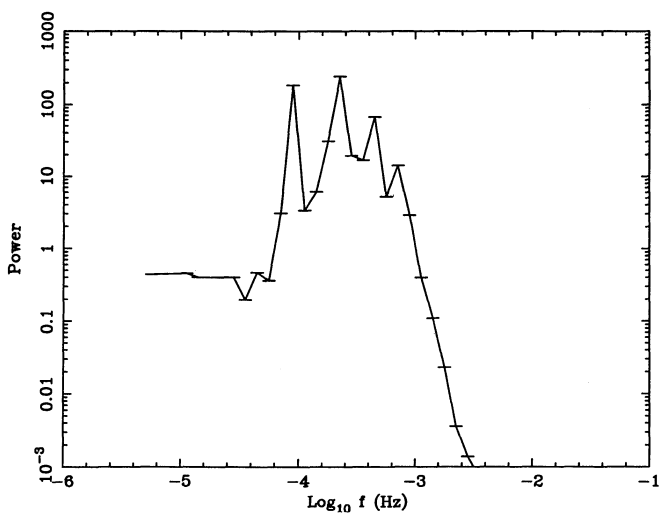

FIG. 8*a*

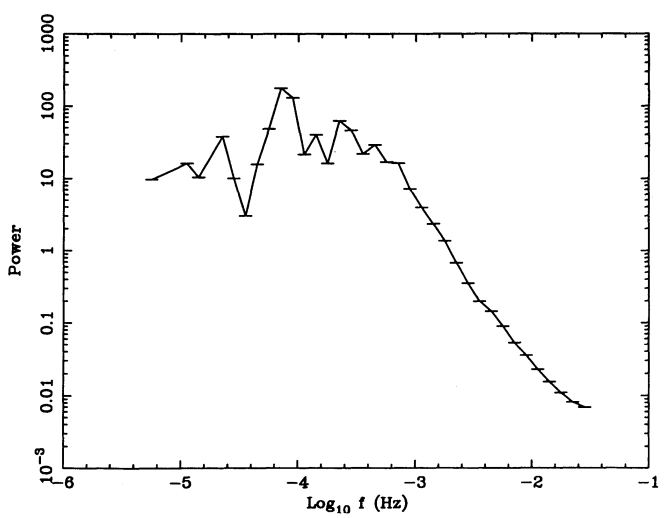

FIG. 8*b*

FIG. 8.—(*a*) Intrinsic power spectrum of the periodic component, calculated from the spline fit to the folded light curve, extended on an even grid to cover the same time span as the real data. (*b*) Resulting power spectrum after this periodic light curve is unevenly sampled in the same way as the *Ginga* data.

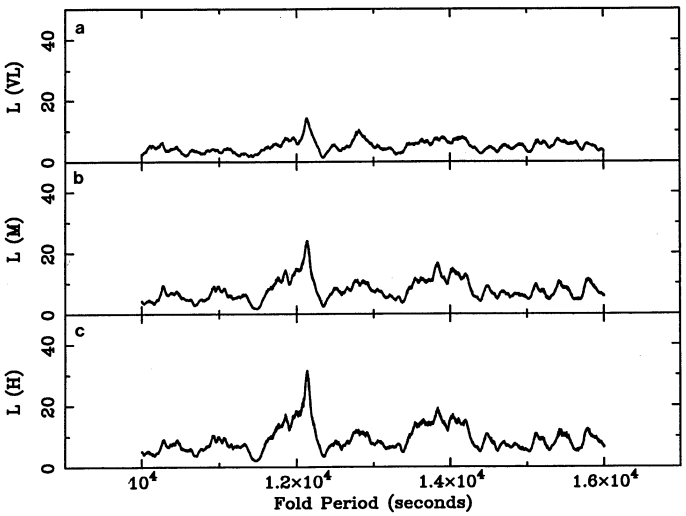

FIG. 9.—(*a–c*) *L*-statistics for the VL, M, and H bands, showing that the periodic component is present in all energy bands.

occultation events. An occulter cannot have infinitely sharp edges, so even if the center is very optically thick there will be a decrease in optical depth toward the edges. As photoelectric opacity decreases with energy, this means that the edges of the material are optically thick at low energies but not at higher energies, so the dip occurs first at low energies and then progressively at higher energies. Similarly, when the occulter uncovers the source, it becomes optically thin first at high energies and then at low, leading to the progressively narrower dip with increasing energy.

For occultation, the amplitude of the dip should be the same in each energy band if the same fraction of the source is covered with optically thick ($\tau \gg 1$) material. Hence the fact that residual flux is seen in the dip in the L, M, and H bands means either that the source is marginally larger than the Thomson thick material, or that the material is not completely opaque ($\tau \sim 1$). Both of these situations can explain the greater depth of the dip in the VL band, the first because the optically thick portion of the occulter is larger at low energies and the second because the optical depth during the center of the dip is largest in the VL band, leading to less possibility of leakage. Whatever its origin, the larger amplitude of variability in the VL band than predicted by the L-band spline gives rise to a hardening of the spectrum during the "dips" in the periodic component. Since the periodic component is dominant in the variability, this produces some if not all of the flux-softness correlation seen by Turner et al. (1992) in the total (as opposed to the folded) light curve. Hence the interpretation of this correlation as evidence for a warm (partially ionized) absorber is not secure, because this behavior can also be repoduced from a periodic occultation of the source.

4.4. *Periodic or Quasi-periodic?*

A knowledge of the stability and coherence length is extremely important in determining the origin of the periodicity. Since the data are unevenly sampled, there is no possibility of using the width of the peak in the spectrum to define the coherence over the 3 days of the *Ginga* observation. Splitting the data into three sections (defined by the two large gaps in the light curve; Fig. 1) was also inconclusive, as each of the

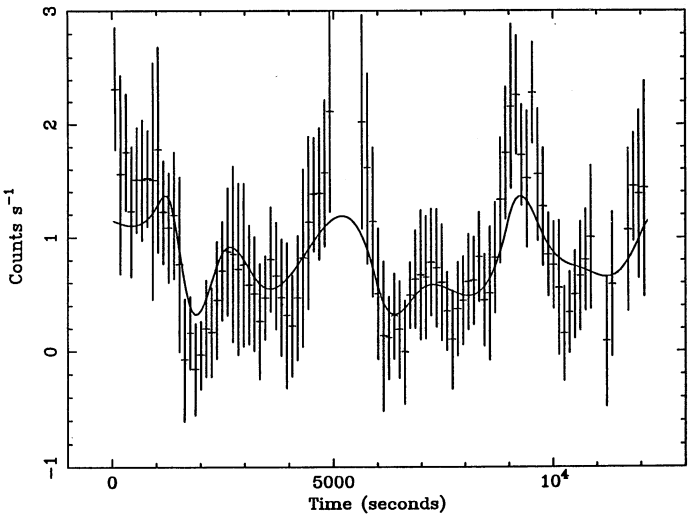

FIG. 10*a*

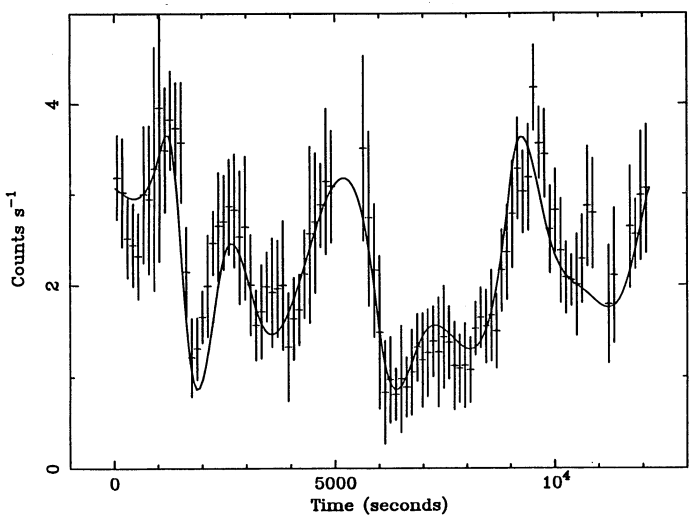

FIG. 10*b*

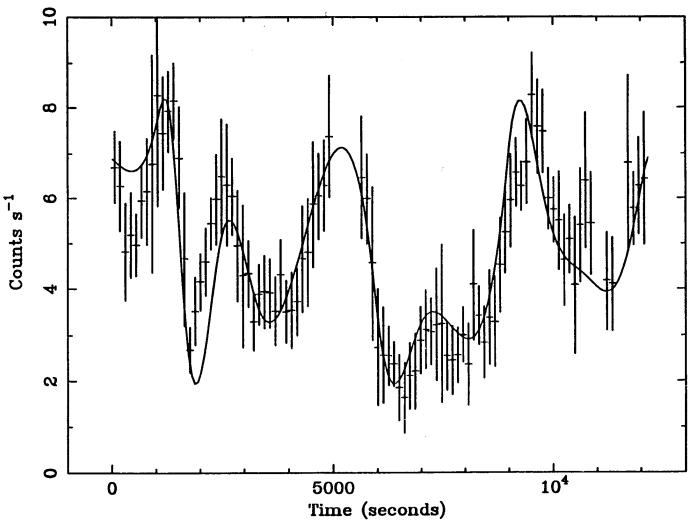

FIG. 10*c*

FIG. 10.—(*a–c*) light curves of the VL, M, and H bands folded on the L-band best-fit period of 12,130.9 s. Superposed is the scaled L-band spline fit, showing a possible though not significant variation of the folded light curve shape with energy.

three sections is then less than a day in total length, giving less than seven periods of a 12,132 s modulation and broadening and weakening the peak in L.

More indirect methods suggest that the signal is periodic over time scales of years. First, the *EXOSAT* data, which are evenly sampled, are consistent with a single period over the 2 years of observations (Fiore et al. 1992b); second, the period derived from *EXOSAT* of 12,200 ± 100 s is consistent with that found here from the *Ginga* data. Also, it is intriguing that the *HEAO 1* A-2 rapid variability was in the form of sharp dips (Tennant & Mushotzky 1983) and that the paper questioning the reality of the variability used as its argument that the events happened at regular intervals (Beall Wood, & Yentis 1986).

Several sharp dips, which are identified with the periodic component, are seen in the second *Ginga* observation of the object (Kunieda et al. 1991). The time of the one sharp dip seen in the observation, in terms of UT day of 1989, is $118.2154 \pm 7.4 \times 10^{-4}$, and those of the three dips seen in the second *Ginga* observation (Kunieda et al. 1991) are $470.8947 \pm 7.4 \times 10^{-4}$, $471.8782 \pm 7.4 \times 10^{-4}$, and $472.0171 \pm 7.4 \times 10^{-4}$. These errors correspond to an estimated accuracy of ±64 s in reading the midpoint of the dip from the expanded light curves in Kunieda et al. (1991). Barycentric corrections can be neglected since the observations are taken almost exactly a year apart, so the relative position of the Earth is the same. Using the period obtained from the folding analysis has the result that the dips in this second observation are separated by 7 and 1 cycles, respectively, and the period for the second observation alone is 12,128 ± 15 s. Taking the dip in the first observation as cycle 1, the cycle number of the first dip in the second observation was varied between 2509 and 2516. A linear fit, which assumes that the signal is completely periodic and phase-coherent between the two observations, is a good description of the data ($\chi^2 = 2.1$ for 3 degrees of freedom) for a best-fit perod of ~12,130 s. Figure 11 shows that there

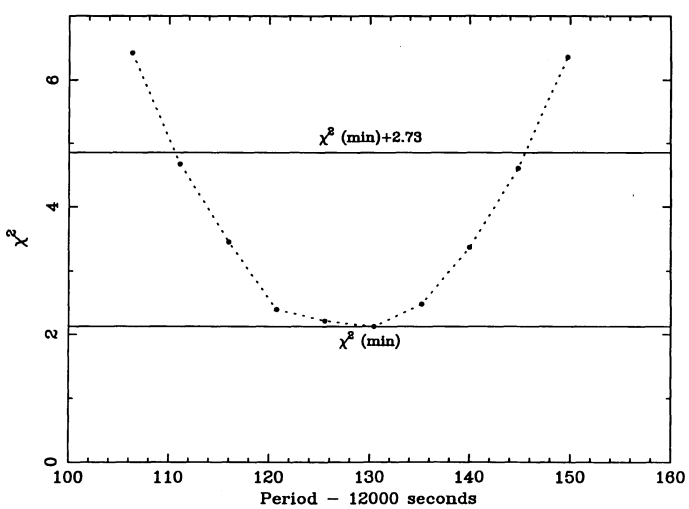

FIG. 11.—The χ^2 for fitting the four dip centroid times (one dip from the data set and three from that of the second *Ginga* observation of the source described by Kunieda et al. 1991) with a linear function. The minimum χ^2 of 2.2 for 3 degrees of freedom shows that the data are consistent with being completely periodic and phase-coherent, though this is not required. However, the period is not completely constrained, as six other solutions besides the period given by the minimum in χ^2 are also acceptable at the 90% confidence limit.

are seven other acceptable periods, ranging from 12,110 to 12,145 s at the $\Delta\chi^2 = 2.73$, i.e., 90%, confidence level. However, the assumption of phase coherence, i.e., the requirement that the dips are at integer numbers of cycles, means that only the periods marked (with their extremely small error bar) are acceptable. Only two of these periods are within 3 σ of the best period from the folded light curve analysis (§ 4.3), namely, 12,130.4 ± 0.05 and 12,135.2 ± 0.05 s. The phenomenon is consistent with being completely periodic and phase-coherent over a time scale of ~1 yr, and in this case the upper limit on the rate of change of period is $\sim 10^{-9}$.

The putative phase coherence must break down over periods of more than a few years, since the *EXOSAT* and *Ginga* folded light curves have very different shapes. Because the two *Ginga* data sets do not *require* phase coherence, the possible change in period alone is given by the difference between the periods derived from the separate data sets, i.e., 18 s in 1 yr of $\dot{P} \leq 5 \times 10^{-7}$, similar to the upper limit on the rate of change of period found from the *EXOSAT* data (Fiore et al. 1992b).

4.5. *Power Spectrum Simulations with Periodic Component*

The periodic component obviously has a large influence on the shape of the power spectrum, though it cannot account for all the observed power (§ 4.3). This leads to the idea of a two-component model, with both broad-band power and a periodic component. The broad-band power is modeled in the same way as described in § 4.2, but its normalization is controlled by the relative normalization of power in the broad-band and periodic components.

The form of the final light curve, $x_s(t)$, is dependent on the type of model assumed for the periodicity. For an occulting screen, the period is made by a modulation of the intrinsic source light curve, and hence obtained by multiplying the light curve from the broad-band power, $x_{bb}(t)$, with that from the periodic component , $x_p(t)$, so that $x_s(t) = x_{bb}(t)x_p(t)$. In contrast, if the period is envisaged as being from an additional X-ray source, its contribution is independent of the intrinsic source light curve, and the final light curve is given by an addition, so that $x_s(t) = x_{bb}(t) + x_p(t)$. These two possibilities are very different in theory (the first has a power spectrum that is the convolution of the power spectrum of the broad-band power with that of the periodic component, whereas the latter is a simple sum) but indistinguishable in practice. We chose the former, multiplicative, method, since the period does not seem to be an additional component in the light curve: the Kunidea et al. (1991) data go to zero intensity.

Since the folded light curve has the same average as the original data, the light curve from the broad-band power is normalized to unity. Its variance, σ_{bb}^2, can be found in terms of the variances of the periodic component, σ_p^2, and desired total variance, σ_s^2, as

$$\frac{\sigma_s^2}{x_s^2} = \frac{\sigma_{bb}^2}{x_{bb}^2} + \frac{\sigma_p^2}{x_p^2} . \quad (6)$$

Because $\sigma/\bar{x}$ is the definition of the rms amplitude, this equation gives the rms amplitude of the broad-band noise as 0.210 (see § 4.3), and so a variance of 4.43×10^{-2}.

Figure 12*a* shows a plot of χ^2 versus the power-law slope, α, for the same coarse grid in α used in § 4.2. The power spectra have 35 frequency bins, but the model uses the average, variance, error bar variance, folded light curve, and periodic fraction as calculated from the real data, so the number of

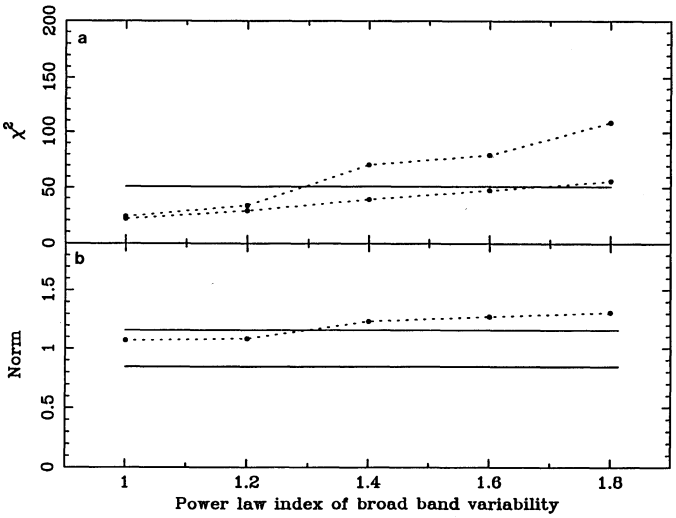

FIG. 12.—(*a*, *b*) As in Figs. 4*a* and 4*b*, but with a periodic component given by the spline fit to the folded light curve as well as a power-law power spectrum to describe the variability. Here there are two solutions, $\alpha = 1.0$ and 1.2, both of which give an acceptable χ^2 and normalization.

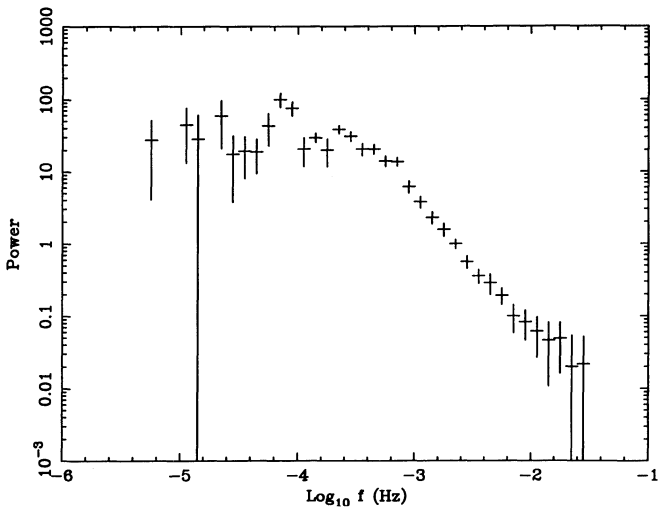

FIG. 14.—Best-fit simulated power spectrum containing the spline-fit periodic component and the broad-band power-law power spectrum with $\alpha = 1.1$. This is a good description of the variability ($\chi^2 = 23.2$ for 30 dof), as shown by a comparison with the power spectrum of the data (Fig. 3).

degrees of freedom is reduced to 30. For the model to be acceptable at the 1% level requires $\chi^2 < 51$. Even with fixed normalization, there are acceptable solutions found for $\alpha = 1.0$ and 1.2, and allowing the normalization of the power spectrum to vary gives a further reduction in χ^2 (see the lower curve in Fig. 12*a*). For this model the factors by which the best-fit normalization differs from the total variance are now within the 15% (i.e., 3 σ error) on the calculated intrinsic variance (see Fig. 12*b*) for $\alpha = 1.0$ and 1.2.

A finer grid in α was then used to get the best fit and error on the power-law index. Figure 13 shows the resulting χ^2, leading to a best-fit of $\alpha = 1.1$. The minimum in χ^2 is complex, but a conservative error range of at least 90% confidence is $0.95 \leq \alpha \leq 1.15$. Figures 14 and 15 show respectively the best-fitting model and its residuals with respect to the observed power spectrum of NGC 6814.

The 10 simulated light curves from the best-fit model were then analyzed using the folding techniques described in § 4.3. The peak in the *L*-statistic was found for each simulation, and this distribution of peak times then gave a mean and error on the mean of 12,132 ± 3 s, showing that there is no systematic offset introduced by the red noise/uneven sampling and giving a true estimate of the error in the period.

4.6. *High-Frequency Break*

An acceptable fit can be obtained by fitting a power law and periodic component to the power spectrum, and the residuals to the best-fit model (Fig. 15) show no indication of a high-frequency cutoff that could indicate a smallest size scale of the emission region. These model smulations were redone with a high-frequency break of $\Delta\alpha = 1$ in the stochastic, power-law variability. This is probably a conservative break, as current

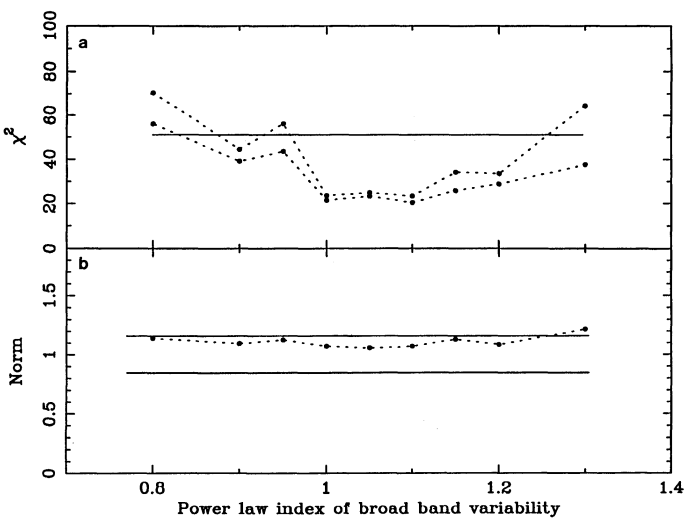

FIG. 13.—As in Figs. 4 and 12, but with a finer grid in α, showing the region of the best fit to the power spectrum. The minimum corresponds to $\alpha = 1.1$.

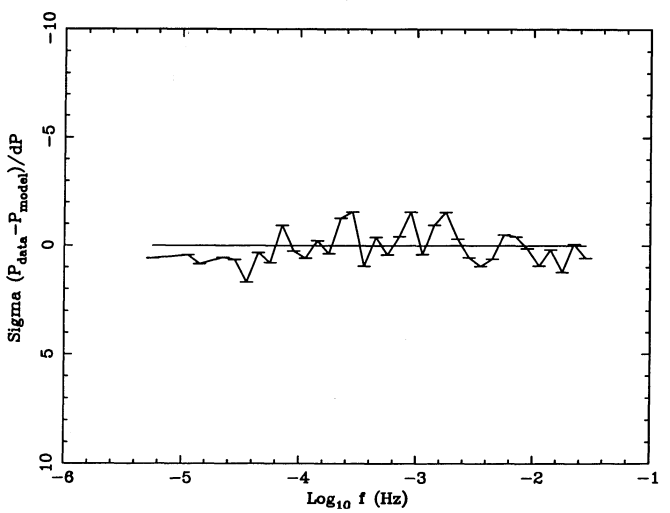

FIG. 15.—Residuals from the fit in Fig. 14

models predict a variety of rollover shapes, from exponential (spots on a disk) to $\Delta\alpha \sim 2$ (shot noise). Three trial break frequencies were used, of 3×10^{-3}, 10^{-2}, and 2×10^{-2} Hz, corresponding to time scales of 300, 100, and 50 s, respectively. The latter two are indistinguishable from an unbroken power law, with $\chi^2 \sim 20$, but the 300 s break has $\chi^2 = 28$, i.e., is a significantly worse fit. The fits were repeated using $P_{\rm HF,noise}$, with its smaller error on the power at high frequencies. Although the χ^2 values increased slightly, as expected, the pattern was still the same, with the data still being unable to distinguish the break at 100 or 50 s from a power law. Hence the high-frequency break must be at higher frequencies than 3×10^{-3} Hz, so the smallest size scale of the emission region is less than 300 lt-s.

5. DISCUSSION

An analysis of the power spectrum and folded light curve of *Ginga* data from NGC 6814 confirms the existence of the periodic component seen in *EXOSAT* data, and the best estimate for the time scale of $12{,}132 \pm 3$ s is consistent with that derived from *EXOSAT* ($12{,}200 \pm 100$ s). This periodic component accounts for at least 75% of the total variability present and has an rms amplitude of 36%. The sharp drops in the light curve which give the most extreme variability are connected with the periodicity. Until the nature of the periodic component is understood, these rapid events cannot be assumed to be a good measure of the intrinsic variability. Calculations of the efficiency or compactness of the source from these events may be misleading.

A second *Ginga* observation of the source taken 1 year later showed a similar periodic dip behavior. Using these dips as markers, the two data sets are consistent with being completely periodic and phase-coherent, giving an upper limit to the rate of change of the period of $\dot{P} < 10^{-9}$. Coherence is not a required property, however, and indeed the *EXOSAT* data are definitely not phase-coherent with the *Ginga* data, since the structure in the folded light curves is different. From the derived periods of the two data sets, the upper limit on the rate of change of period is $\dot{P} \leq 5 \times 10^{-7}$, similar to the limits set by the *EXOSAT* data.

Underlying the periodic component, the power spectrum shows the more usual $f^{\sim -1}$ broad-band power. The data are consistent with this extending down to the shortest time scales, although the noise dominates below time scales of 300 s. Any high-frequency break then has a characteristic time scale of less than 300 s.

5.1. *The Origin of the Periodicity*

5.1.1. *Rotational Time Scale*

The nature of the periodic component is a subject of much debate. Because it is seen over 6 yr between the *EXOSAT* and *Ginga* observations, it must represent some fundamental and stable time scale of the system, which tends to point toward rotation as the origin of the period. As noted by Mittaz & Branduardi-Raymont (1989), if the time scale represents that of the rotation period of the innermost stable orbit of the accretion disk, then the mass of the black hole is $2.6 \times 10^7\ M_\odot$ or $3.8 \times 10^8\ M_\odot$ (corresponding to an inner edge radius of 2.3×10^{13} cm or 5.7×10^{13} cm) is a Schwarzschild or extreme Kerr geometry, respectively.

There are problems with such an identification. The periodic component contains some extremely rapid variability, on time scales 500 times shorter than this putative orbital time scale. This requires that the source be much smaller than the orbital radius, but it does not address the central issue as to what actually causes the modulation observed. An intriguing way around this problem is given by Abramowicz et al. (1992), who postulate a highly inclined system, with the X-rays emitted from a "hot spot" on the disk. This gives the period as the rotational time at the appropriate disk radius, and the rapid time variability in the periodicity comes from the gravitational lensing of the spot at it passes behind the black hole. An ensemble of smaller, short-lived spots then give the more usual AGN f^{-1} broad-band noise power. Since there is a continual accretion of shots on the viscous time scale of $t_{\rm visc} = \alpha^{-1}(H/R)^{-2}t_{\rm dyn}$, this requires that $\alpha(H/R)^2 \ll 6.4 \times 10^{-5}$ for the period to be observed for more than 6 yr. For an inclination of 85° this gives $\alpha \ll 8.4 \times 10^{-3}$, a not unreasonable limit, although the requirement that the inner regions of the disk are observed at such high inclination angles restricts the luminosity to $L/L_{\rm Edd} \sim H/R \leq 0.08$. However, the predicted iron-line width is rather greater than that derived from current observations (Abramowicz et al. 1992), and while the complex, phase-dependent line profile may allow some discrepancy, the model is also unable to explain trivially the very different structure of the folded light curve between the *EXOSAT* and *Ginga* epochs.

5.1.2. *Disk Instabilities*

Pair instabilities (White & Lightman 1990; Moskalik & Sikora 1986) in the inner disk are also discussed in Mittaz & Branduardi-Raymont (1989). A high central temperature leads to a limit-cycle behavior: pair production decreases the temperature, switching off the pair production, which leads to an increase in the central temperature, which switches on pair production. The period expected is of the right order, with $P \sim 4 \times 10^{-4}\ M/M_\odot$ s, but large changes in disk structure are expected, as it makes the transition from hot and extended to cooler and thin, which have a strong spectral signature. This is in conflict with the results that the spectral index is constant over changing flux state (Kunidea et al. 1990). Similarly, accretion disk instabilities involve large changes in the disk structure and so are improbable because of the constancy of the spectral index with flux. Other disk time scales such as the hydrodynamic, thermal, or viscous time scales could be made to work with an appropriate choice of both the disk viscosity and the radius at which the periodic component arises, but this seems a very ad hoc approach.

5.1.3. *QPO Mechanisms*

An analogy with the quasi-periodic oscillations (QPOs) seen in the galactic X-ray binaries (see Lewin, van Paradijs, and van der Klis 1988 for a review) seems initially to be a strong candidate. There are three spectral branches on which QPOs are detected—the horizontal, normal, and flaring branches—where the order is thought to represent increasing accretion rate. Beating between the magnetosphere of the magnetically controlled accretion flow and the rotation of the neutron star is the best description of the origin of the QPO on the horizontal branch. This is thought to give way to the normal branch when the accretion rate becomes too large to be constrained by the magnetic field. The resulting quasi-spherical inward accretion then interacts with the outwardly flowing radiation, and the balance between accretion-producing radiation which inhibits further accretion forms a quasi-periodic limit-cycle behavior.

As the accretion rate increases still further, the flow becomes super-Eddington, with bipolar outflow. However, none of these seem to be applicable to NGC 6814, the first because the surface of the neutron star is a key ingredient (and black holes have no hard surface), the second because the rms variability in the predicted light curve is small, and the third because the periodic behavior has been seen over a long period of time, implying a stability incompatible with super-Eddington systems.

5.1.4. *Occultation Structures*

A more promising line of argument is that the periodic time scale is indeed a rotational time scale, but that the geometry is such that there is another "special" radius apart from the inner radius, e.g., as in a thick disk. Is such a geometry, where the observer is looking in a direction close to the opening angle of the disk funnel, small changes in the height of the disk edge could cause periodic occultation. It is this argument that is used to explain the strange QPO behavior of the galactic black hole candidate LMC X-1 (Ebisawa, Mitsuda, & Inoue 1989). While optical images show NGC 6814 to be a system viewed close to face-on, the angular momentum in the central regions may be decoupled from the underlying galaxy, especially if the AGN activity is fueled by a merger or interaction which can introduce a different angular momentum structure.

Occultation is also suggested by the spectral fits of Kunieda et al. (1990), who split the data into four intensity states and found that the absorbing column is greater in the lowest intensity spectra. As most of the variability is connected with the periodic component, the four spectra are almost equivalent to taking four intensity cuts in the periodic component. Hence there is more absorbing material in the dips than in the "flares." Also, the spectral index in these four states is constant, making it unlikely that some periodic flare is superposed on a separate continuum source.

The time scales required are consistent with the occulter being identified with the disk. The shape of the disk is stable on the viscous time scale $t_{\rm visc} = \alpha^{-1}(H/R)^{-2}t_{\rm dyn}$, where α is the disk viscosity parameter, which must be very much less than unity for stability (Abramowicz, Calvani, & Nobili 1980; Madau 1988). Small perturbations on the edge of the disk are only stable on time scales on the order of the hydrodynamic time scale. As this is similar to the dynamical time scale, the detailed shape of the disk edge is unlikely to be maintained over many rotations of the disk. This predicts that while the period will remain stable, phase coherence will be lost on observations over long time scales as the edge structure changes.

Another way of obtaining structure on the disk is given by Siemiginowska & Elvis (1992), who show that there is a change in vertical structure at the hydrogen recombination radius. The abrupt increase in opacity traps the radiation, so the disk scale height increases. Again, structure on the peak occults the X-ray source in the same way as the previous thick-disk model, although the recombination radius is typically too large to explain the period observed here.

Occultation by nondisk structures has been suggested by Syer, Clarke, & Rees (1991) and Sikora & Begelman (1992). These both use a star captured by the black hole to occult the source, with the plane of the stellar orbit inclined to the plane of the disk. Such an orbit will precess, causing a change in the observed structure of the periodic component.

If occultation by whatever means is the origin of the periodic component, then the X-ray source must be less than 300 lt-s (the time taken for the sharp drop) in projected size. If the iron line comes from reprocessing in an extended region, the size scale is expected to be larger than the continuum size scale. Because the iron line contributes $\sim$20% of the total counts in the M band (for a 300 eV equivalent width line), this predicts that either the dip time scale should be longer (if all of the iron-line–producing region is occulted) or that the dip is smaller in relative depth than in the L band (if the line-producing region is not occulted, so that the nonperiodic line fluxes dilutes the dip). Instead, the data show that the light curve across the iron-line band is consistent with the same amount of variability and the same short-dip time scale, showing that the iron line is predominantly produced within a region smaller than 300 lt-s across (confirming the result of Kunieda et al. 1990) and is affected by the periodic modulation. This leads to ambiguity in interpreting the asymmetry of the line and continuum cross-correlation (Done et al. 1992). Naively, the effect suggests an intrinsic small (50–200 s) line lag, but such an effect can also be produced by an external occulter. This will be dealt with in more detail in a later paper.

5.2. *Additional Information from the Spectrum*

While occultation by structure on the disk or by a captured star may be able to explain the existence of the periodic component, we caution that the situation may be more complex, since it seems likely that the X-ray emission is anisotropic. The limit on the size scale from the lack of high-frequency break gives a lower limit to the compactness of the source, irrespective of the maximum dL/dt measured from the light curve which is contaminated by the periodic component. Assuming that all the X-ray luminosity is from the minimum size, then $l = L\sigma_T/(Rm_e c^3) > 30$. Any γ-rays produced by the source are likely to be reprocessed through photon-photon pair production, as the X-ray source is optically thick to this process. However, then the flat spectrum becomes difficult to interpret. A spectral index flatter than $\alpha = 0.5$ requires that the source be photon-starved (Done & Fabian 1989; Zdziarski & Coppi 1991). However, pair production tends to steepen the spectrum of even photon-starved sources. Done & Fabian (1989) show that this spectral index can only be obtained for $l_s/l_h \sim 0.01$ if pair production is important. A periodic component intrinsic to the radiation mechanism can actually be ruled out on these grounds since in this range of parameters, large luminosity changes produce large spectral index changes (Done & Fabian 1989) which are not seen in the data (Kunieda et al. 1990).

It is difficult to reconcile NGC 6814 with being photon-starved. The large iron-line equivalent width shows that there are large amounts of material illuminated by the X-ray source, and the rapid variability of the line implies that this material lies very close to the source. X-rays intercepted by the material which are not reprocessed as iron-line or reflection continua are thermalized in the source. As only $\sim$10% of photons are reprocessed, this provides copious soft photons. These will be intercepted by the X-ray source, and so violate the requirement that the source be photon-starved. Some degree of anisotropy seems then to be the only way to explain the emission seen.

6. SUMMARY

The existence of the periodic component seen in NGC 6814 with *EXOSAT* at 12,200 $\pm$ 100 s is confirmed by a power

spectrum and folded light curve analysis of unevenly sampled *Ginga* data. A comparison of the power spectra produced from simulated light curves with that observed enables the intrinsic shape of the power spectrum of the source to be determined despite the distortions introduced by the window function. The best estimate for the period is 12,132 ± 3 s, where the error is that derived from simulations. A second *Ginga* observation of this source is consistent with being completely periodic and phase-coherent with these data, for seven possible periods (each determined to within 50 ms), spanning a range from 12,110 to 12,145 s. This implies an upper limit to the rate of change of period of $\sim 10^{-9}$ if the light curves are truly phase-coherent but as this is not required by the data, the conservative upper limit is $\dot{P} \leq 5 \times 10^{-7}$. Phase coherence is definitely not maintained over time scales longer than ~1 yr, since the shape of the periodic component in the light curve differs strongly between the *Ginga* and *EXOSAT* observations.

In addition to the periodic component, with rms amplitude of 36%, there is broad-band power which can be modeled as a power law such that $P(f) \propto f^{-1.1}$. This continues unbroken down to time scales of at least 300 s, where the measurement noise dominates. Apart from the strong periodicity, the power spectrum of NGC 6814 is indistinguishable from that of other low-luminosity AGNs.

The large amount of power in the periodic component, and its stability, both suggest occultation of the source as its origin. There is tentative support for this in the energy dependence of the sharp-dip events found in the folded light curve analysis and in a more detailed study of the sharp-dip events in the second *Ginga* observation by Leighly et al. (1992). There are several possible models for the occulter, either a star orbiting the black hole or vertical structure on an accreton disk. Both these models suffer from being strongly fine-tuned. However, the more physically appealing model of a highly inclined system with gravitational lensing of X-ray emission regions on the disk may have difficulty in accounting for the differing shape of the periodicity seen in the *Ginga* and *EXOSAT* data.

Possibly all these models fail to give a convincing explanation for all the properties of the variability because of their assumption that the source is intrinsically isotropic. The contradictory requirements that the source be photon-starved in order to see a flat spectrum, and yet that there be large amounts of material close to the source that reprocess the X-rays in order to see the large equivalent width in the iron line, suggest that the X-rays are anisotropic in some respect.

Many people have helped with their comments, but we would especially like to thank Fabrizio Fiore for many enlightening discussions about the meaning of power spectra; Rich Kelly, Koji Mukai, and Alan Smale for their knowledge of folding techniques; Steve Shore for his understanding of how these are affected by uneven sampling; Marek Abramowicz for discussions of the "hot spots" model; and Rick Edelson for motivation. C. D. acknowledges receipt of a National Research Council–GSFC Research Associateship.

APPENDIX A

Midlayer data are excluded from this analysis, as are data from "contact" orbits passing through the South Atlantic Anomaly, because of their higher background. The selection criteria used are SUD (surplus above upper discriminator) ≤ 7.8 counts s^{-1} per detector, COR (geomagnetic cutoff rigidity) ≥ 9 GeV/c, Earth elevation angle greater or equal to 5°, pointing direction of the satellite within 0°.1 of the source position, and omission of the first and last bin after any detector configuration changes. These data were then "cleaned" by rejecting any points more than 3 σ away from a linear relationship between SUD and PI-MONI (particle monitor), and the background subtraction carried out from a nearby and temporally close observation.

A detailed discussion of the effects of data selection and background subtraction for faint sources can be found in Fiore et al. (1992c). The data in channels 38–47 (~25–35 keV) are used to test for any possible systematic effects of the high-energy background subtraction. An extrapolation of the mean spectrum to this energy band gives an expected count rate of ~0.08 counts s^{-1} from the source, whereas fitting the data directly gives 0.087 ± 0.02 counts s^{-1} (1 σ error bars). However, this mean count rate is not consistent with being constant, with $\chi^2 = 2124$ for 1847 degrees of freedom. The predicted variance from the small source count rate in the energy band is $\sim 1.7 \times 10^{-3}$, assuming that it varies in the same manner as the L band, whereas the measured variance is $\sim 5.2 \times 10^{-2}$. Thus the error bars in this very high energy band are underestimated on average by ~0.2 counts s^{-1}.

This additional variance is not contributing to the periodicity. Figure 16*a* shows the fold statistic (L) for this very high energy band light curve. The excess variance and can be seen from the fact that the mean L over the band is 1.5 rather than 1, but there is no structure at ~12,100 s, showing that the peak seen in source-dominated data (Figs. 6 and 9) is not from any background subtraction error. This can also be seen from the fact that a different data selection (top-layer data only, but South Atlantic Anomaly included, SUD ≤ 8 counts s^{-1} per detector, 10 GeV/$c \leq$ COR ≤ 20 GeV/c, attitude restricted to pointing direction within 0°.4 from the source, 6° $\leq$ Earth angle $\leq$120°, cleaned by removing points more than 3 σ from a linear SUD–PI-MONI relationship) and background subtraction method (that of Hayashida et al. 1989) shows the same strong signal present at ~12,100 s (Fig. 16*b*). That the different uneven sampling pattern from this analysis gives the peak in L at 12,133 s, consistent with that derived in § 4.5, also confirms that the peak cannot be due to an alias of the data windowing.

To estimate the effect of the additional variance in the H band, the rms residuals after background subtraction shown in Fiore et al. (1992c) are summed over the relevant channels. This gives the maximum systematic error, since it assumes that the residuals have the same sign in each channel, and is of order ~0.2 counts s^{-1} in both the very high energy band used above and the H band. Thus a similar excess variance is predicted in the H band due to systematic errors associated with background subtraction, but the large observed variance (see Table 1) means that this is at most a 2% effect.

While these data are identical to those used in the discussion of the spectrum in Turner et al. (1992), channels 7–9 are not excluded from this analysis. The possible problem of contamination from a radioactive decay line, while of great importance in a spectral

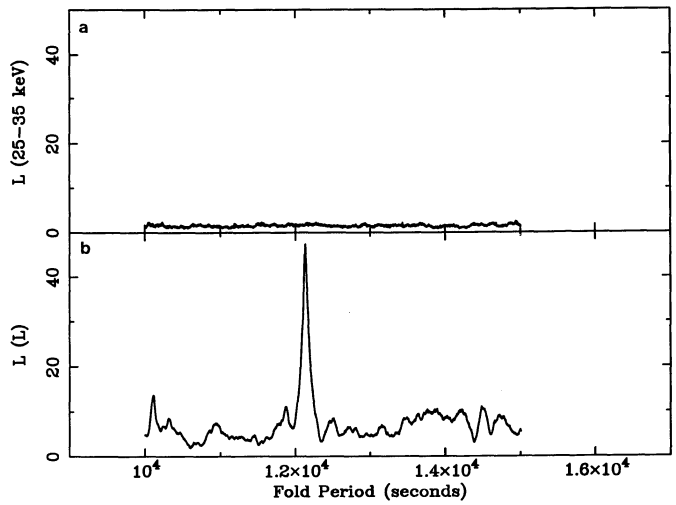

FIG. 16.—(*a*) *L*-statistic of the very high energy band light curve, 25–35 keV. The source signal is very weak, so these channels are dominated by background. While there is some additional variance from systematic errrors in the background subtraction, this does not contribute to the periodicity. (*b*) *L*-statistic from the L-band light curve of the data extracted with different selection criteria (so giving a different sampling pattern) and different background subtraction. The peak occurs at 12,133 s, consistent with the previous results, showing that the period is neither from background subtraction effects nor from aliasing.

analysis, especially in connection with the detailed shape of the red wing of the iron line, has negligible effect in timing analysis, as can be seen from the fact that the different (Hayashida) background subtraction gives the same L-band significance of the periodic component (Fig. 16*b*).

APPENDIX B

In Fourier analysis the light curve is broken into cosine components of different frequency, and for evenly sampled data the squares of the amplitudes of these components give the intrinsic power spectrum, $P(f)$. Reversing this, if the form of the intrinsic power spectrum is given, then the evenly sampled light curve can be described by a sum of cosines of amplitude $\propto[P(f)]^{1/2}$. A random distribution (Press et al. 1986) of phase, $\phi(f)$, with frequency is assumed, and then the simulated light curve, $x_s(t)$, at time t is given by the inverse Fourier series, summed over all frequencies

$$x_s(t) \propto \sum_{f=1/T}^{1-\Delta t} \frac{\sin \pi f \Delta t}{\pi f \Delta t} [P(f)]^{1/2} \cos [2\pi f t - \phi(f)] , \tag{B1}$$

where the term $(\sin \pi f\Delta t)/(\pi f\Delta t)$ is a correction to account for the suppression of high-frequency noise through the binning process (van der Klis 1989). The light curve is produced so that it spans the entire time of the observed data set, evenly sampled on the shortest sample time of the real data. Then, in order to mimic the effects of the *Ginga* window, the light curve is resampled on the same uneven pattern as the data by linearly interpolating between the closest points to the times at which the observed data points were taken.

In order to normalize the simulated light curve to match the data, the simulated light was scaled so that its average was equal to that of the original data and its variance was equal to the error-corrected estimate of the intrinsic variance, $\sigma^2_{\rm int}$, of the variability, $\sigma^2_{\rm int} = \sigma^2_{\rm tot} - \sigma^2_{\rm err}$ (see § 4.1). An error bar, $dx_s(t)$, was then added to the data, chosen as a Gaussian-distributed random variable with variance given by the linear relationship between the count rate and its error bar, $dx(i)$ and $x(i)$. A linear approximation is valid, since the count rate in *Ginga* is dominated by the background rather than by the source. For the L band, a good description of the error is $dx(i) = 2.79 \times 10^{-2}x(i) + 0.955$.

This gives a simulated data set with an intrinsic power spectrum $P(f)$ (chosen here to be of power-law form) and with the same mean, variance, error bar, binning, and sampling characteristics as the original data set. The power spectrum of this light curve can then be calculated as for the original data.

REFERENCES

Abramowicz, M. A., Calvani, M., M., & Nobili, L. 1980, ApJ, 242, 772
Abramowicz, M. A., Bao, G., Fiore, F., Lanza, A., Massaro, E. Perola, G. C., Spiegel, E. A., & Szuszkiewicz, E. 1992, in The Physics of Active Galactic Nuclei, ed. S. J. Wagner & W. J. Duschl (Heidelberg: Springer-Verlag), in press
Abramowitz, M., & Stegun, I. A. 1972, Handbook of Mathematical Functions (New York: Dover)
Beall, J. H., Woods, K. S., & Yentis, D. J. 1986, ApJ, 308, 563
Davies, S. R. 1990, MNRAS, 244, 93
———. 1991, MNRAS, 251, 64P
Deeming, T. J. 1975, Ap&SS, 36, 137
Done, C., & Fabian, A. C. 1989, MNRAS, 240, 81
Done, C., Madejski, G. M., Mushotzky, R. F., Turner, T. J., Koyama, K., & Kunidea, H. 1992, in Frontiers of X-Ray Astronomy, ed. Y. Tanaka & K. Koyama (Tokyo: Universal Academy Press)
Ebisawa, K., Mitsuda, K., & Inoue, H. 1989, PASJ, 41, 519
Edelson, R. A., & Krolik, J. H. 1988, ApJ, 333, 646
Fiore, F., Done, C., Edelson, R., & Angelini, L. 1992a, in preparation
Fiore, F., Massaro, E., & Barone, P. 1992b, A&A, in press
Fiore, F., Massaro, E., Perola, G. C., & Piro, L. 1989, ApJ, 347, 171
Fiore, F., Perola, G. C., Matsuoka, M., Yamauchi, M., & Piro, L. 1992c, A&A, in press

Guilbert, P. W., & Rees, M. J. 1988, MNRAS, 233, 475
Hayashida, K., et al. 1989, PASJ, 41, 373
Hoel, P. G. 1984, An Introduction to Mathematical Statistics (New York: Wiley)
Kunieda, H., Turner, T. J., Awaki, H., Koyama, K., Mushotzky, R. F., & Tsusaka, Y. 1990, Nature 345, 786
Kunieda, H., et al. 1991, in Iron Line Diagnostics in X-Ray Sources, ed. A. Treves, G. C. Perola, & L. Stella (Berlin: Springer-Verlag), 241
Lawrence, A., Pounds, K. A., Watson, M. G., & Elvis, M. 1987, Nature 325, 694
Leahy, D. A., Darbro, W., Elsner, R. F., Weisskopf, M. C., Sutherland, P. G., Kahn, S., & Grindlay, J. E. 1983, ApJ, 266, 160
Leighly, K., Kunieda, H., & Tsuruta, S. 1992, in Testing the AGN Paradigm, ed. S. S. Holt, S. G. Neff, & C. M. Urry (New York: AIP), 93
Lewin, W. H. G., van Paradijs J., & van der Klis M. 1988, Space Sci. Rev., 46, 273
Lightman, A. P., & White, T. R. 1988, ApJ, 335, 57
Lomb, N. R. 1976, Ap&SS, 39, 447
Madau, P. 1988, ApJ, 327, 116
McHardy, I. M. 1990, in Proc. 23d ESLAB Symp. (ESA SP-296; Paris: ESA), 2, 1111
McHardy, I. M., & Czerny, B. 1987, Nature, 325, 696
Mittaz, J. P. D., & Branduardi-Raymont, G. 1989, MNRAS, 238, 1029
Moskalik, P., & Sikora, M. 1986, Nature 319, 649
Pounds, K. A., Nandra, K. P., Stewart, G. C., George, I. M., & Fabian, A. C. 1990, Nature, 344, 132
Press, W. H. 1978, Comm. Astrophys., 7, 103
Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1986, Numerical Recipes (Cambridge: Cambridge Univ. Press)
Press, W. H., & Teukolsky, S. A. 1988, Comput. Phys., November December, 77
Scargle, 1982, ApJ, 263, 835
Schwarzenberg-Czerny, A. 1989, MNRAS, 241, 153
Siemiginowska, A., & Elvis, M. 1992, ApJ, submitted
Sikora, M., & Begelman, M. C. 1992, Nature, 356, 224
Syer, D., Clarke, C. J., & Rees, M. J. 1991, MNRAS, 250, 505
Tennant, A. F., & Mushotzky, R. F. 1983, ApJ, 264, 92
Turner, M., et al. 1989, PASJ, 41, 345
Turner, T. J., Done, C., Kunieda, H., Mushotzky, R. F., & Awaki, H. 1991, in IAU Colloq. 129, Structure and Emission Properties of Accretion Disks, ed. C. Bertout, S. Collin, J.-P. Lasota, & J. Tran Thanh Van (Gif-sur-Yvette: Editions Frontières), 547
Turner, T. J., Done, C., Mushotzky, R. F., Madejski, G. M., & Kunieda, H. 1992, ApJ, 391, 102
yan der Klis, M. 1989, ARA&A, 27, 517
White, T. R., & Lightman, A. P. 1990, ApJ, 352, 495
Zdziarski, A. A., & Coppi, P. S. 1991, ApJ, 376, 480
Zdziarski, A. A., Ghisellini, G., George, I. M., Svensson, R., Fabian, A. C., & Done, C. 1990, ApJ, 363, L1