High Speed Photometer
Pulsar Timing and Light Curve Reduction

Jeffrey W Percival

Space Astronomy Laboratory,

University of Wisconsin,

Madison, WI 53706

May 19, 1992

Abstract

This memo describes in detail the reduction of the 5 High Speed Photometer observations of the Crab Pulsar. This analysis shows that the correlation of the Hubble Space Telescope clock with UTC is well within the design specification of 10 milliseconds, and may be within 1 millisecond.

1 Introduction

The Crab pulsar was observed on five separate occasions by the High Speed Photometer (HSP) on the Hubble Space Telescope (HST). In each case the sample time was $11/1.024 \times 10^6 \approx 10.74$ microseconds. Table 1 gives the circumstances of each observation. Table 2 gives the histograms of the raw data.

Table 1

Observation	Date	Filter	Duration (s)
v0qy0102	1991 Oct 17 16:09:03.60	F400LP	2160
v0qy0201	1991 Oct 18 06:39:02.61	F400LP	1800
v0qy0301	1991 Oct 19 06:49:52.61	F400LP	1800
v0qy0401	1991 Oct 20 07:00:43.61	F400LP	1800
v0ui0103	1992 Jan 21 07:58:06.67	F160LP	1680

Table 1: Summary of HSP observations of the Crab pulsar

The light curves were produced with the following procedure. First, the definitive HST ephemeris was obtained from the Flight Dynamics Facility at Goddard Space Flight

Table 2

	Counts per sample							
Observation	6	5	4	3	2	1	0	total
v0qy0102	0	0	12	329	13908	733379	200180372	200928000
v0qy0201	1	0	9	384	17002	844570	166577394	167439360
v0qy0301	0	0	5	254	11550	655336	166772215	167439360
v0qy0401	0	0	2	243	10321	601035	166827759	167439360
v0ui0103	0	0	0	0	164	72859	156203457	156276480

Table 2: Histograms of the raw data.

Center. This ephemeris gives the 6-element HST barycentric J2000 state vector at one minute intervals, expressed as the Modified Julian Date in units of 100 nanoseconds. For each of these state vectors, the barycentric state vector of the earth was obtained from the JPL DE-200 Planetary Ephemeris. The sum of these vectors gives the barycentric state vector $\mathbf{v}$ of the HST, and the arrival time correction is then $\mathbf{v} \cdot \hat{\mathbf{p}}/c$ where $\hat{\mathbf{p}}$ is the unit vector in the direction of the Crab Pulsar and c is the speed of light. Cubic spline interpolation in the resulting table of corrections was used to arrive at barycentric corrections within each minute. The phase of each non-zero sample was calculated using the procedure outlined in the Jodrell Bank Monthly Ephemeris (Lyne & Pritchard 1992). Figures 1- 5 show the resulting mean light curves.

The following sections describe each of these steps in detail.

2 Position of HST

The definitive HST ephemeris is given in files that ultimately originate in the Goddard Space Flight Center's Flight Dynamics Facility. We obtained the appropriate files from the Science Institute's Data Management Facility. Table 3 gives the files used, and their time coverage in Modified Julian Dates. The accuracy of the HST state vector can be estimated from the data files. Consecutive files share one common time point. The last time point of the previous file is the same as the first time point of the next file. The state vectors for these time points, however, differ according to their respective regression solutions. Comparing

pbag0000r.orx with pbai0000r.orx, we find two state vectors for MJD 48547.0: (6632.978, -1360.692, -1706.646) km and (6632.962, -1360.749, -1706.670) km. The difference is (-0.016, 0.057, 0.024) km, or a total displacement of 64 meters. The light travel time for this distance is less than 2% of our sample time, and is negligible for our purposes.

Table 3

Ephemeris file	Start (MJD)	Stop (MJD)
pbag0000r.orx	48545.0	48547.0
pbai0000r.orx	48547.0	48549.0
pbak0000r.orx	48549.0	48551.0
pc1k0000r.orx	48641.0	48643.0

Table 3: Definitive Ephemeris files used for HST position

3 Time of day

Events on the HST are driven by a master 1.024 MHz oscillator. This same oscillator also drives the HSP microprocessor and data collection hardware. The HST Science Data Formatter (SDF) keeps a 32-bit software clock with a resolution of 125 ms. The HSP software clock is also a 32-bit counter, but the resolution is 1 ms. Assigning absolute times to the photometry samples collected by the HSP requires establishing a time correlation for each of these clocks. That is, one must calibrate the zero point of the HSP clock with respect to the HST clock, and the zero point of the HST clock with respect to UTC in order to map HSP clock values onto absolute time.

The HSP/HST clock correlation is established by logging the HST clock value at the instant that the HSP clock is set to zero.

The HST/UTC correlation is established by time tagging the receipt of certain telemetry signals at the White Sands ground station. The time tags are then corrected for signal propagation times from White Sands to the TDRS relay satellite in use, and from the TDRS to the HST. The largest remaining uncertainty is in the HST telemetry hardware. Depending on the telemetry format in use at the time of the calibration, the uncertainty can vary from 1 ms (format A) to 8 ms (format C). The calibration is performed daily, and a regression is

performed against the accumulated data to extract the coefficients of a quadratic polynomial mapping the HST clock onto UTC. The clock rate (seconds / count) and clock drift (seconds / count^2) are important outputs of this regression.

The absolute time of an HST event is given by

$$T = T_0 + (V_1 - V_0)r_0 + (V_1 - V_0)^2 d_0/2 \quad (1)$$

where T_0 is the absolute time corresponding to vehicle time V_0, V_1 is the vehicle time of the desired event, r_0 is the HST clock rate in seconds per count, and d_0 is the HST clock drift rate in seconds per count squared.

We obtained the results of several regressions from Morlock and Kimmer (1992) and chose a solution, shown in Table 4, whose epoch was close in time to the final Pulsar observation.

Table 4

parameter	value
T_0	1992.004 05:44:00.443
V_0	428038096
r_0	0.1250000009901
d_0	$8.861552441864 \times 10^{-19}$

Table 4: HST/UTC Clock correlation

4 Time of observation

Now, with these two clock correlations, we can convert any given HST or HSP time to absolute time. The one remaining problem is assigning the correct HST time to the first photometry sample. There are two ways to approach this, and in order to appreciate the distinction between them we must describe the sequence of events that leads up to the collection of science data with the HSP. In the photometry mode used for the Crab Pulsar, the HSP is programmed to start integrating the first sample at some prearranged HSP time. That is, when the HSP clock reaches some target value, then exactly 1 ms + 1 HSP clock tick later, the first sample begins (Werner, 1992). Although the HSP collects its samples

in a continuous stream, the HST SDF quantizes the downlink into data packets. For the Crab observations, each packet contains 1920 1-byte photometry samples. When the HSP finishes the 1920th sample, it begins an interaction with the SDF to read out the packet. The HSP asserts a signal called Frame Start, and then sometime later, typically 1-2 ms, the SDF asserts a signal called Line Start. The data are then read out of the HSP. Before shipping the packet out to the downlink, the SDF inserts the HST clock value at which the Line Start was generated into the packet header (Rankin, 1992). Although the HST clock has a resolution of only 125 ms, the SDF adds an additional 10 bits of resolution, improving it to $125/2^{10}$ ms, or about $122\mu s$.

We can now estimate the HST time of the first photometry sample in two ways: adding the 1 ms + 1 HSP clock tick offset to the programmed start time, or by taking the time embedded in the first data packet and subtracting the time it took to collect the first packet's data, including all the known delays and overheads. We call these two methods the forward method and the backward method. The two principle overheads to account for in the backward method are the amount of time that elapses *after* the 1920th sample has been collected but *before* the HSP asserts the Frame Start, and the delay in the SDF before asserting the Line Start. The HSP delay has been estimated to be 1.068-2.61 ms (Werner, 1992), and the SDF delay is 1-2 ms. Table 5 shows the HST times for the Crab observations using each method. The packet time is the actual value recovered from the first data packet. The "backward time" column has 1920 sample times removed, *but not the HSP or SDF overheads.* The "forward time" column gives the programmed start time, *without the 1 ms + 1 HSP clock tick delay.* We see that the two methods are discrepant by 0.041703 ± 0.003340 HST clock ticks, or about 5.2 ± 0.4 ms. The sense is that the forward method times are *earlier* than the backward method times. Some of the mean discrepancy is due to the neglect of the various delays and overheads. They would move the forward method later by about a millisecond, and would move the backward method earlier by about 3.3 ms (assuming average values for the ranges given above). There remains about 0.9 ms of unknown error.

Table 5

Observation	packet time	backward time	forward time	difference
v0qy0102	373733321.875977	373733321.710977	373733321.672008	0.038969
v0qy0201	374150913.881836	374150913.716836	374150913.672008	0.044828
v0qy0301	374847313.875977	374847313.710977	374847313.672008	0.038969
v0qy0401	375543721.876953	375543721.711953	375543721.672008	0.039945
v0ui0103	439852865.882813	439852865.717812	439852865.672008	0.045804
mean				0.041703
std. dev.				0.003340

Table 5: Timing data (values are in HST clock ticks, 125 ms)

5 Calculation of phase

The phase of each non-zero sample was calculated using the following procedure, given in Lyne & Pritchard (1992). Denoting the pulsar period with P, and letting ν_0 and $\dot{\nu}$ be the pulsar's frequency and frequency derivative at some epoch, we have

$$P = 1/\nu_0, \tag{2}$$

$$\dot{P} = -\dot{\nu}/\nu_0^2, \tag{3}$$

and

$$\ddot{\nu} = 2\dot{P}^2/P^3. \tag{4}$$

The phase at the instant of observation T is then

$$\Phi = \nu_0 T + \dot{\nu}T^2/2 + \ddot{\nu}T^3/6. \tag{5}$$

Table 6 gives the radio epoch, and the phase of the first sample of each observation.

6 Mean light curves

Figures 1-5 show each observation coadded into 3125 phase bins covering one pulsar period. 3125 was chosen to make each phase bin approximately as wide as the $10.74\mu s$ sample width. The backward method times in Table 5 were used to produce the light curves in these figures. Tables 7 and 8 give the phase error of the main peak and the phase drift

Table 6

Observation	Radio Epoch	ν_0 (Hz)	$\dot{\nu}(10^{-10}/s^2)$
v0qy0102	48544 + 0.030786 s	29.9436079304	-3.7748018
v0qy0201	48544 + 0.030786 s	29.9436079304	-3.7748018
v0qy0301	48544 + 0.030786 s	29.9436079304	-3.7748018
v0qy0401	48544 + 0.030786 s	29.9436079304	-3.7748018
v0ui0103	48636 + 0.020498 s	29.9406077550	-3.7737775

Table 6: Radio ephemerides used in the data reduction

during the observation. The phase error and phase drift ($\Delta P/P$) was estimated by dividing the observation into about 200 time slices, and coadding the light curve separately for each slice. The phase of the maximum value was found for each slice, and a least squares fit was performed using the resulting few hundred pairs of points. Figures 6-10 show the peak phase as a function of pulse number, with the linear regression superimposed. The large variance in the data is primarily due to the crude estimate of the peak phase in each of the slices: we chose the phase bin containing the most counts, rather than using a more elaborate technique such as smoothing with curve fitting. Each time slice lasted for only about 10 seconds, and contained just a few hundred counts. The "maximum value" choice would be expected to suffer using the relatively noisy 10-second light curves, but should not do any systematic damage to the regression.

Table 7

Observation	Φ_0	Phase error	Phase drift ($\Delta P/P$)
v0qy0102	6922707.609	-9.66670e-02	-2.507667e-08
v0qy0201	8485862.431	-9.50542e-02	-4.144763e-08
v0qy0301	11092663.576	-9.56528e-02	5.811730e-09
v0qy0401	13699489.437	-9.42972e-02	-2.386960e-08
v0ui0103	16391899.795	-9.58908e-02	-3.832335e-08
mean		-9.55124e-02	-2.458110e-08
std. dev.		8.92000e-04	1.869362e-08
RMS			2.97286e-08

Table 7: Phase errors and drifts (Forward method).

We take the y-intercept of the linear fit to be the phase error with respect to the radio

Table 8

Observation	Φ_0	Phase error	Phase drift ($\Delta P/P$)
v0qy0102	6922707.755	4.823839e-02	-1.618972e-08
v0qy0201	8485862.599	7.154610e-02	-3.709999e-08
v0qy0301	11092663.722	5.051263e-02	-6.236897e-08
v0qy0401	13699489.587	5.380815e-02	4.463694e-12
v0ui0103	16391899.966	7.258340e-02	-4.114319e-09
mean		5.933773e-02	-2.395371e-08
std. dev.		1.179137e-02	2.586712e-08
RMS			3.33026e-08

Table 8: Phase errors and drifts (Backward method).

ephemeris, and the slope to be the phase drift within a given observation. Using an average pulsar frequency of 29.94 Hz to convert the phase errors into time errors, we end up with -3.190 ± 0.02979 ms for the forward method, and 1.982 ± 0.3938 ms for the backward method. Now recall the neglected delays and overheads. The forward method should be increased by about a millisecond, to -2.2 ms, and the backward method should be decreased by about 3.3 ms, to -1.3 ms. The residual discrepancy of about 0.9 ms is what we expected, and is due to the unexplained discrepancy in the time of the first datum. Adding the backward method timing uncertainties in quadrature (using 0.771 ms for the HSP delay and 0.5 ms for the SDF delay) we get a total uncertainty in the calculated arrival times of 1.0 ms. The forward method delays of 1 ms + 1 HSP clock tick are exact, so the total uncertainty is just the dispersion of phase errors, or 0.02979 ms.

We see that each of these methods has an advantage. The forward method gives a larger phase error, but produces a very small variance. We believe that the phase error is due to some small, but highly deterministic, overhead in the HSP, perhaps occurring when the HSP resets its internal clock. In fact, a small phase shift between the HST and HSP clocks could introduce a millisecond delay in the reset (Werner, 1992), thereby shifting the HSP time axis by the amount required to eliminate the discrepancy. Note that the backward method times would not be affected by such a phase shift, because they are derived from the HST clock in the packet headers.

The backward method gives a better match to the radio ephemeris, but with a larger dispersion. We think the real phase error of the backward method may be smaller than the -1.3 ms quoted above, because an average value of 1.5 ms was used for the SDF delay, but with only one active instrument, the SDF delay may be closer to the 1 ms lower limit, and this would bring the backward method phase error to within the 1 ms total uncertainty in the delays and overheads.

These small phase errors and dispersions are a little surprising, given the variability of the HST clock drift term d_0. A plot in Morlock and Kimmer (1992) shows d_0 plotted as a function of time in 1991, and shows rapid and frequent excursions of at least 3×10^{-18} seconds per count^2. Extrapolating such a change from the January solution back to the October Crab observations (90 days, at 8 HST ticks per second), we get a temporal uncertainty of about 6 ms. The January Crab observation fell much closer to the correlation epoch, and so is far less sensitive to the uncertainty in d_0. It is not clear to what extent our small errors and dispersions and due to good luck. A different d_0 would have thrown the October light curves further from the radio ephemeris, increasing the mean error, and further from the January light curve, increasing the dispersion of phase errors. In any case, the HST clock seems to be calibrated to within the 10 ms specification even after folding in all methods, errors, delays, and overheads.

The phase drift provides an estimate of the amount of smearing in the final light curve for each observation. The RMS phase drift of $\Delta P/P = -3 \times 10^{-8}$ over approximately half an hour implies a smearing of about $54\mu s$, or about 5 samples. Referring to Figures 6-10, it is clear that the measured drifts are much smaller than the variance of the data, and are therefore statistically negligible. Figures 11-15 show enlargements centered on the peaks. Note that the peak v0qy0401 shows a profile about as wide as the others, yet its drift measurement in Table 8 is 4 orders of magnitude lower than the others. This too indicates the statistical meaninglessness of the residual phase drift, and bolsters the feeling that the flat tops of the peaks are intrinsic to the pulsar and not an artifact of numerical smearing in the data reduction.

REFERENCES

Lyne, A. G. and Pritchard, R. S. 1992, *Jodrell N Bank Crab Pulsar Timing Results Monthly Ephemeris (March 10, 1992)*

Morlock, S. and Kimmer, E. 1992, *memo to J. Hodges dated January 27, 1992*

Rankin, A. 1992, *Private communication*

Werner, M. L. 1992, *Private communication*

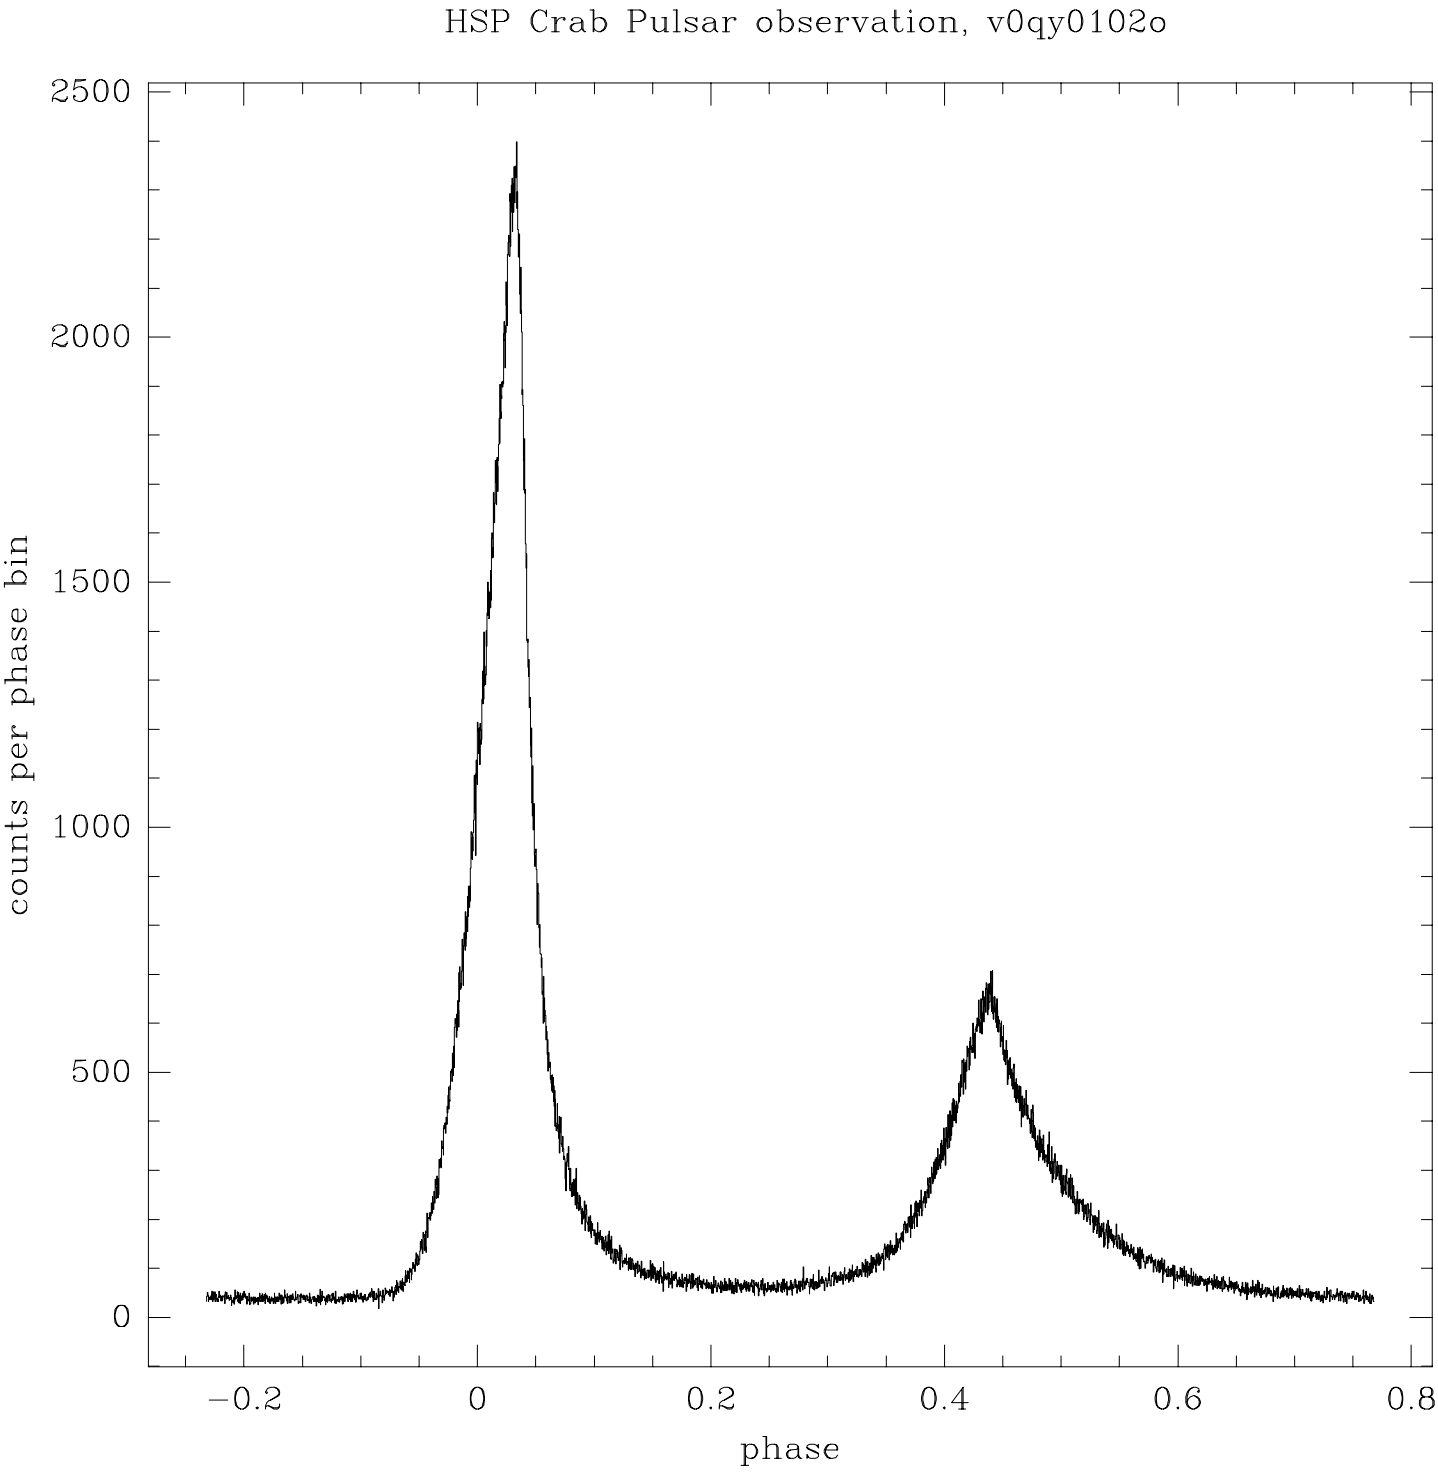
HSP Crab Pulsar observation, v0qy0102o
2500
2000
1500
1000
500
0
counts per phase bin
−0.2
0
0.2
0.4
0.6
0.8
phase

Figure 1:

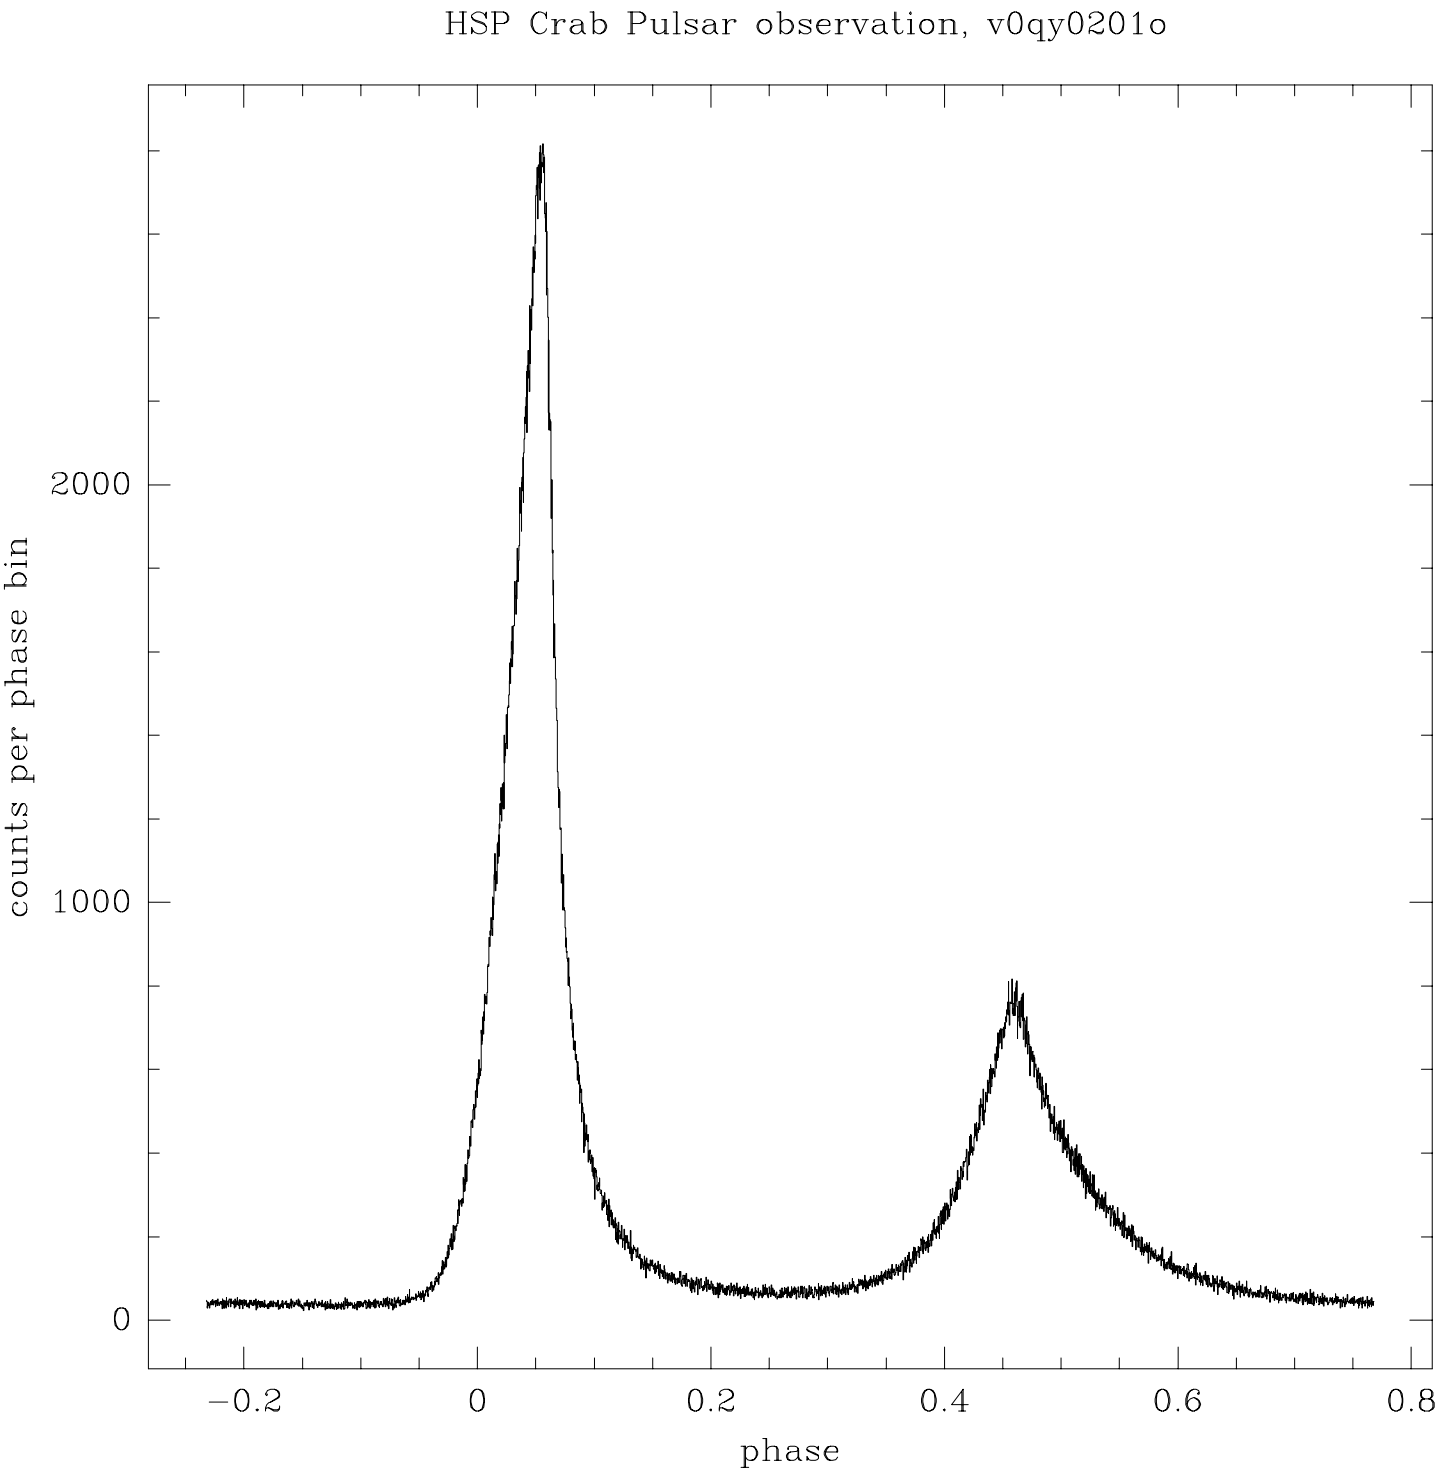
HSP Crab Pulsar observation, v0qy0201o
counts per phase bin
2000
1000
0
−0.2
0
0.2
0.4
0.6
0.8
phase

Figure 2:

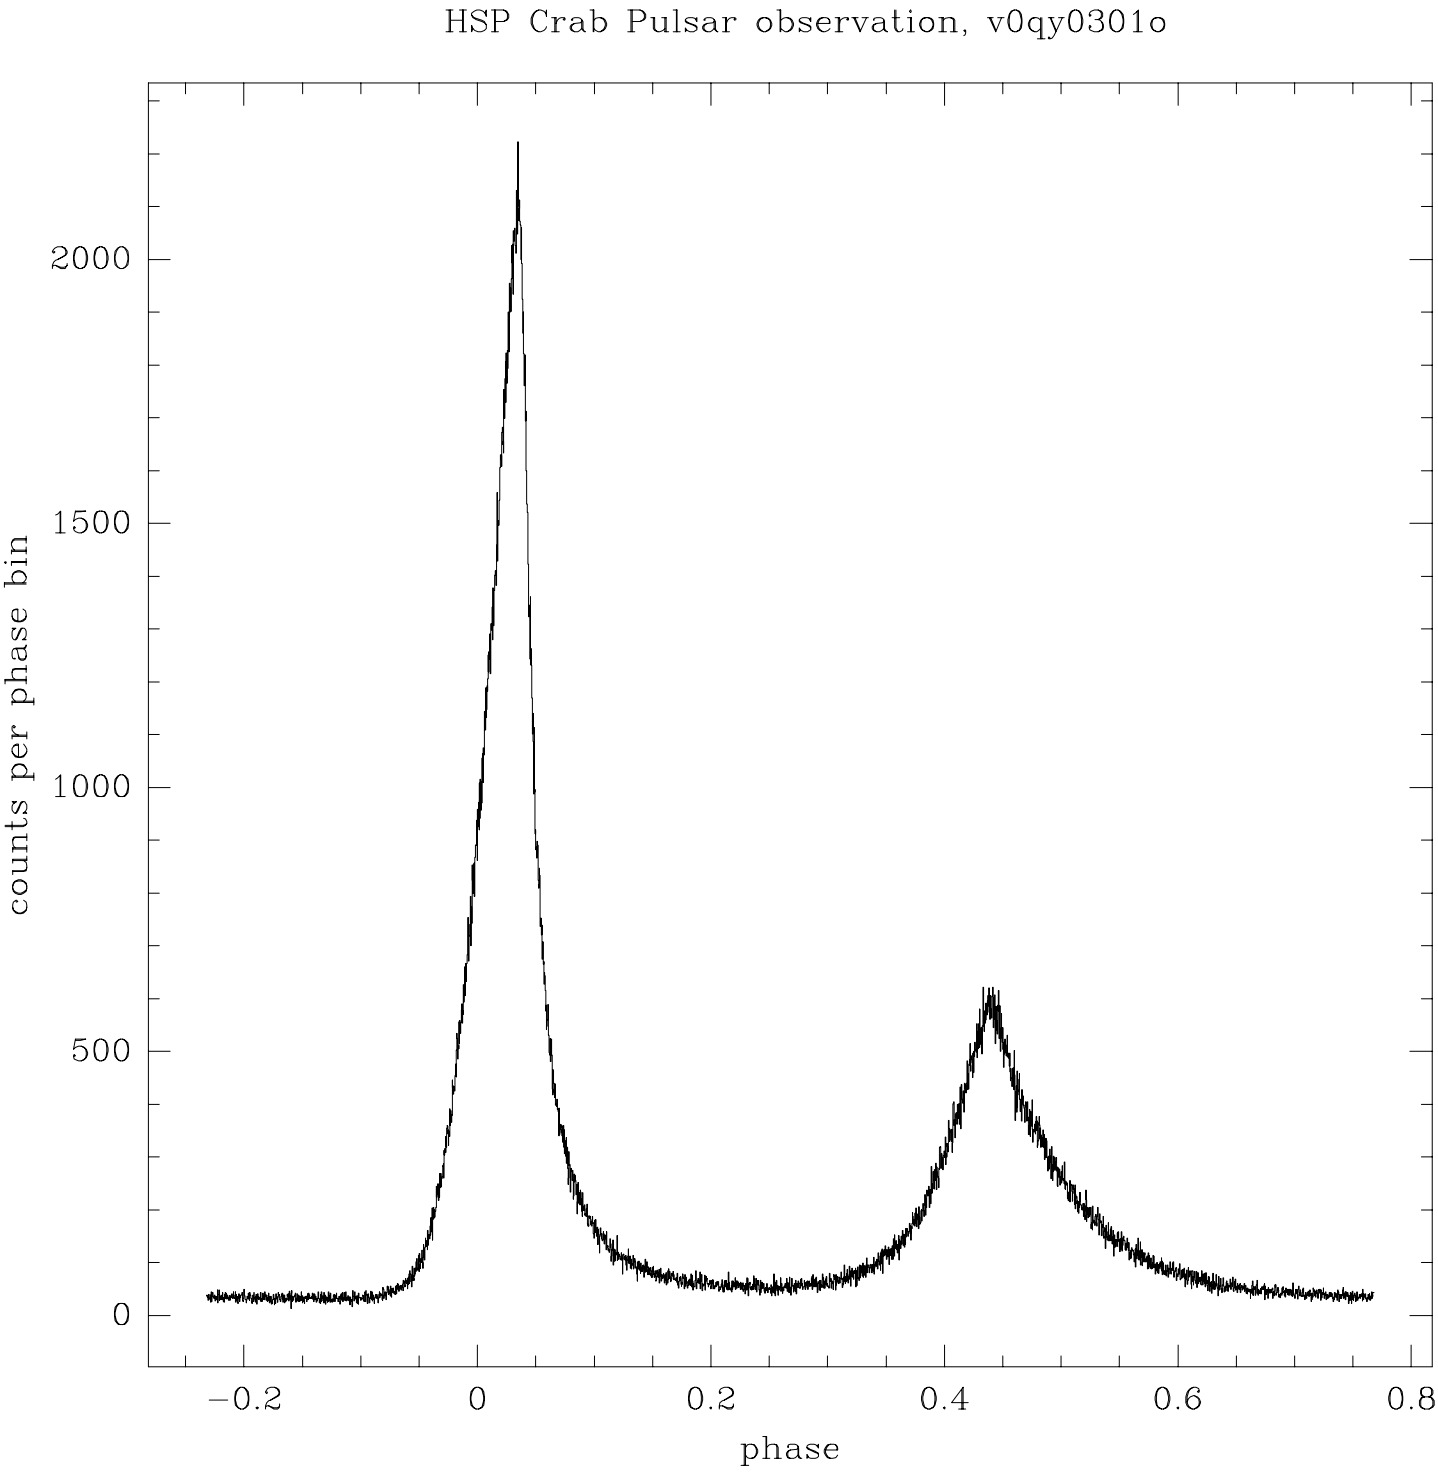
HSP Crab Pulsar observation, v0qy0301o
counts per phase bin
2000
1500
1000
500
0
−0.2
0
0.2
0.4
0.6
0.8
phase

Figure 3:

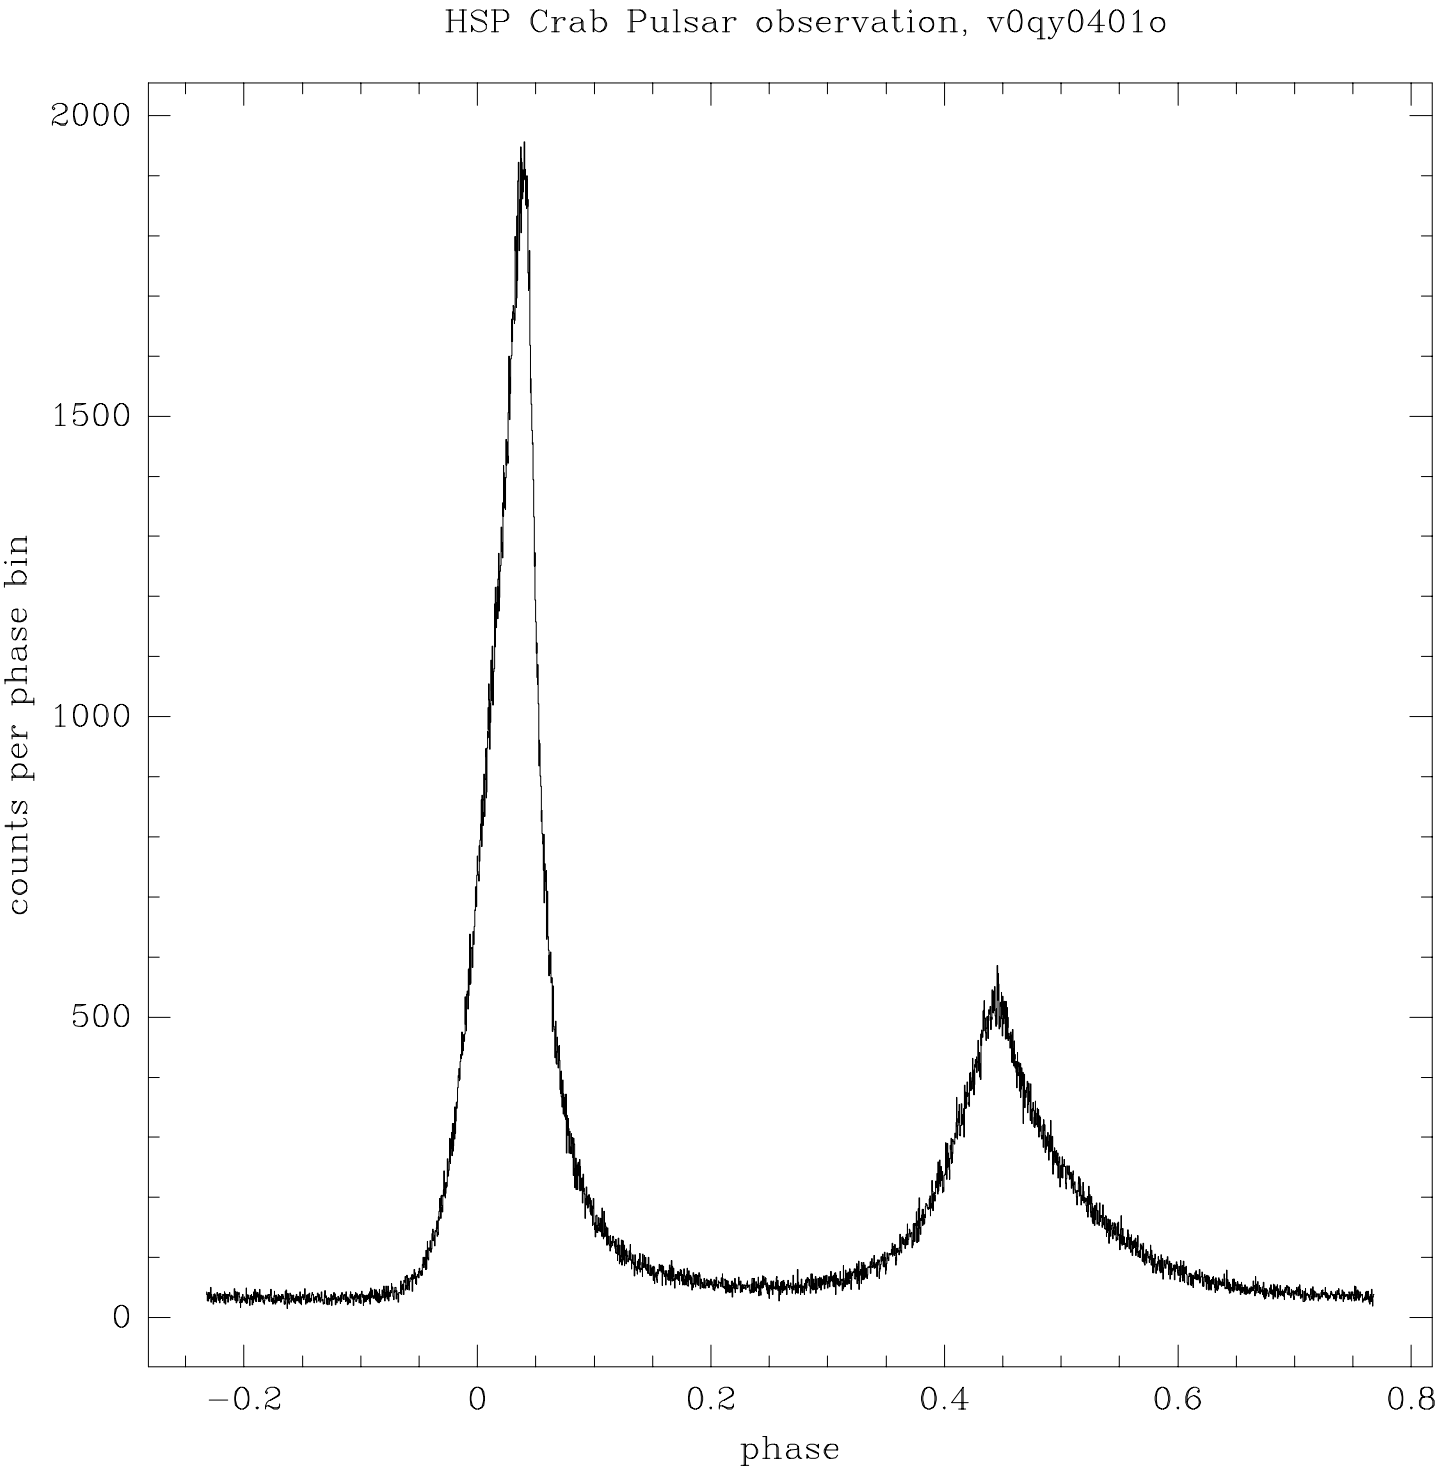
HSP Crab Pulsar observation, v0qy0401o
2000
1500
1000
500
0
counts per phase bin
−0.2
0
0.2
0.4
0.6
0.8
phase

Figure 4:

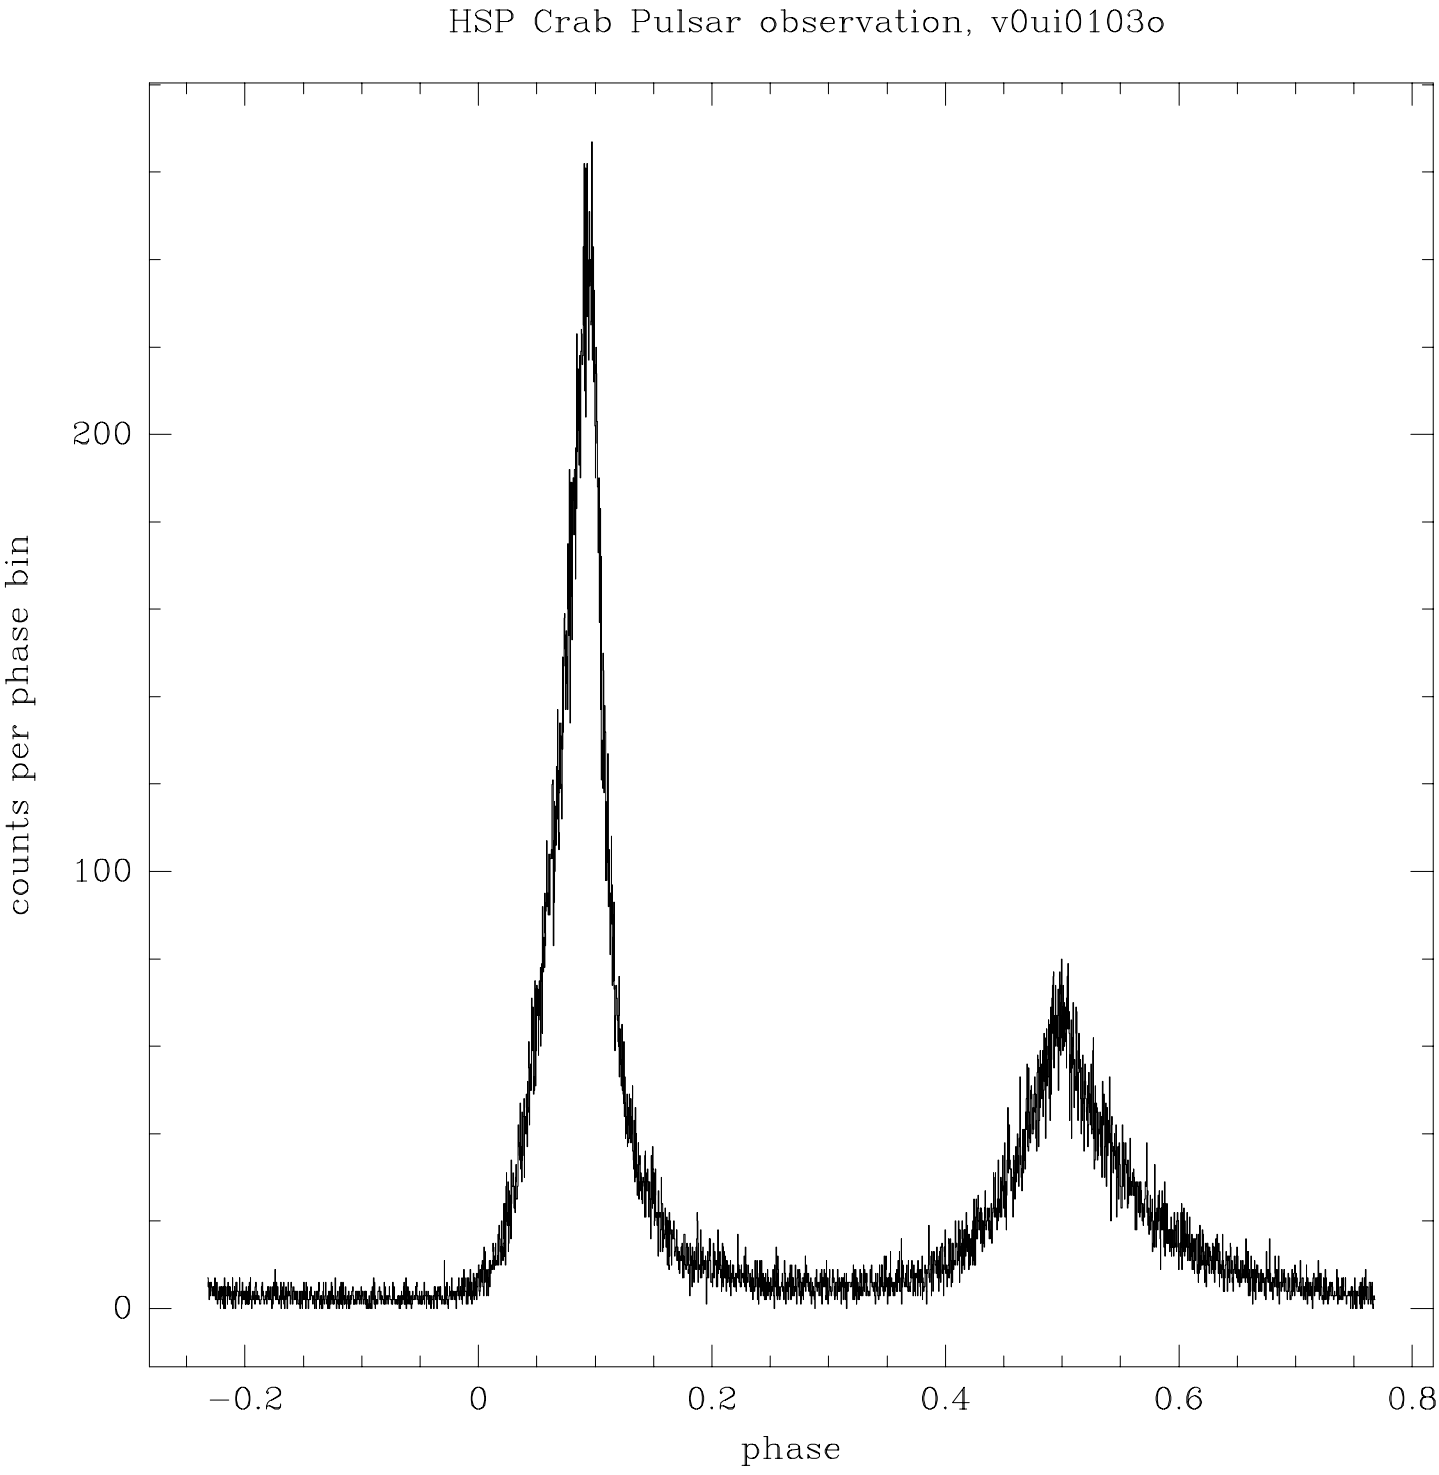
HSP Crab Pulsar observation, v0ui0103o
counts per phase bin
200
100
0
−0.2
0
0.2
0.4
0.6
0.8
phase

Figure 5:

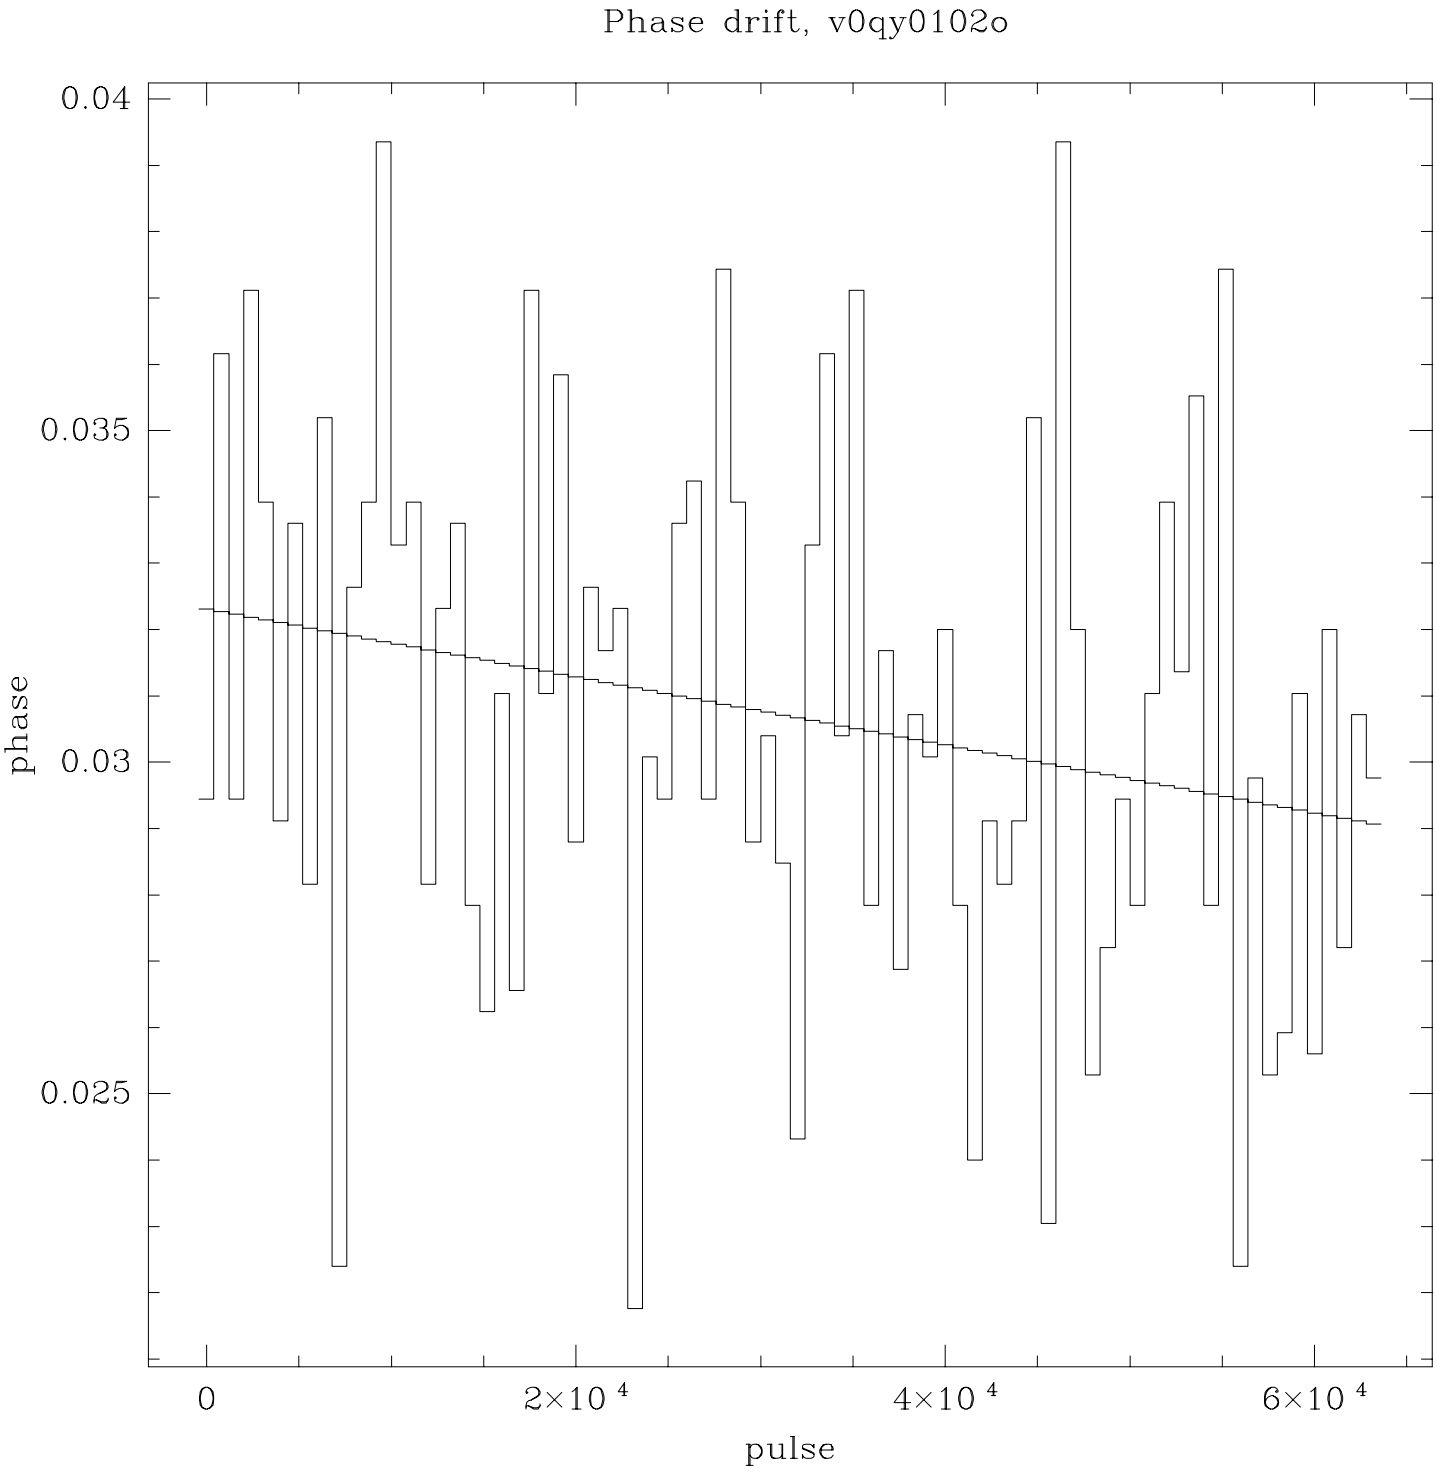
Phase drift, v0qy0102o
0.04
0.035
0.03
0.025
phase
0
2×10 4
4×10 4
6×10 4
pulse

Figure 6:

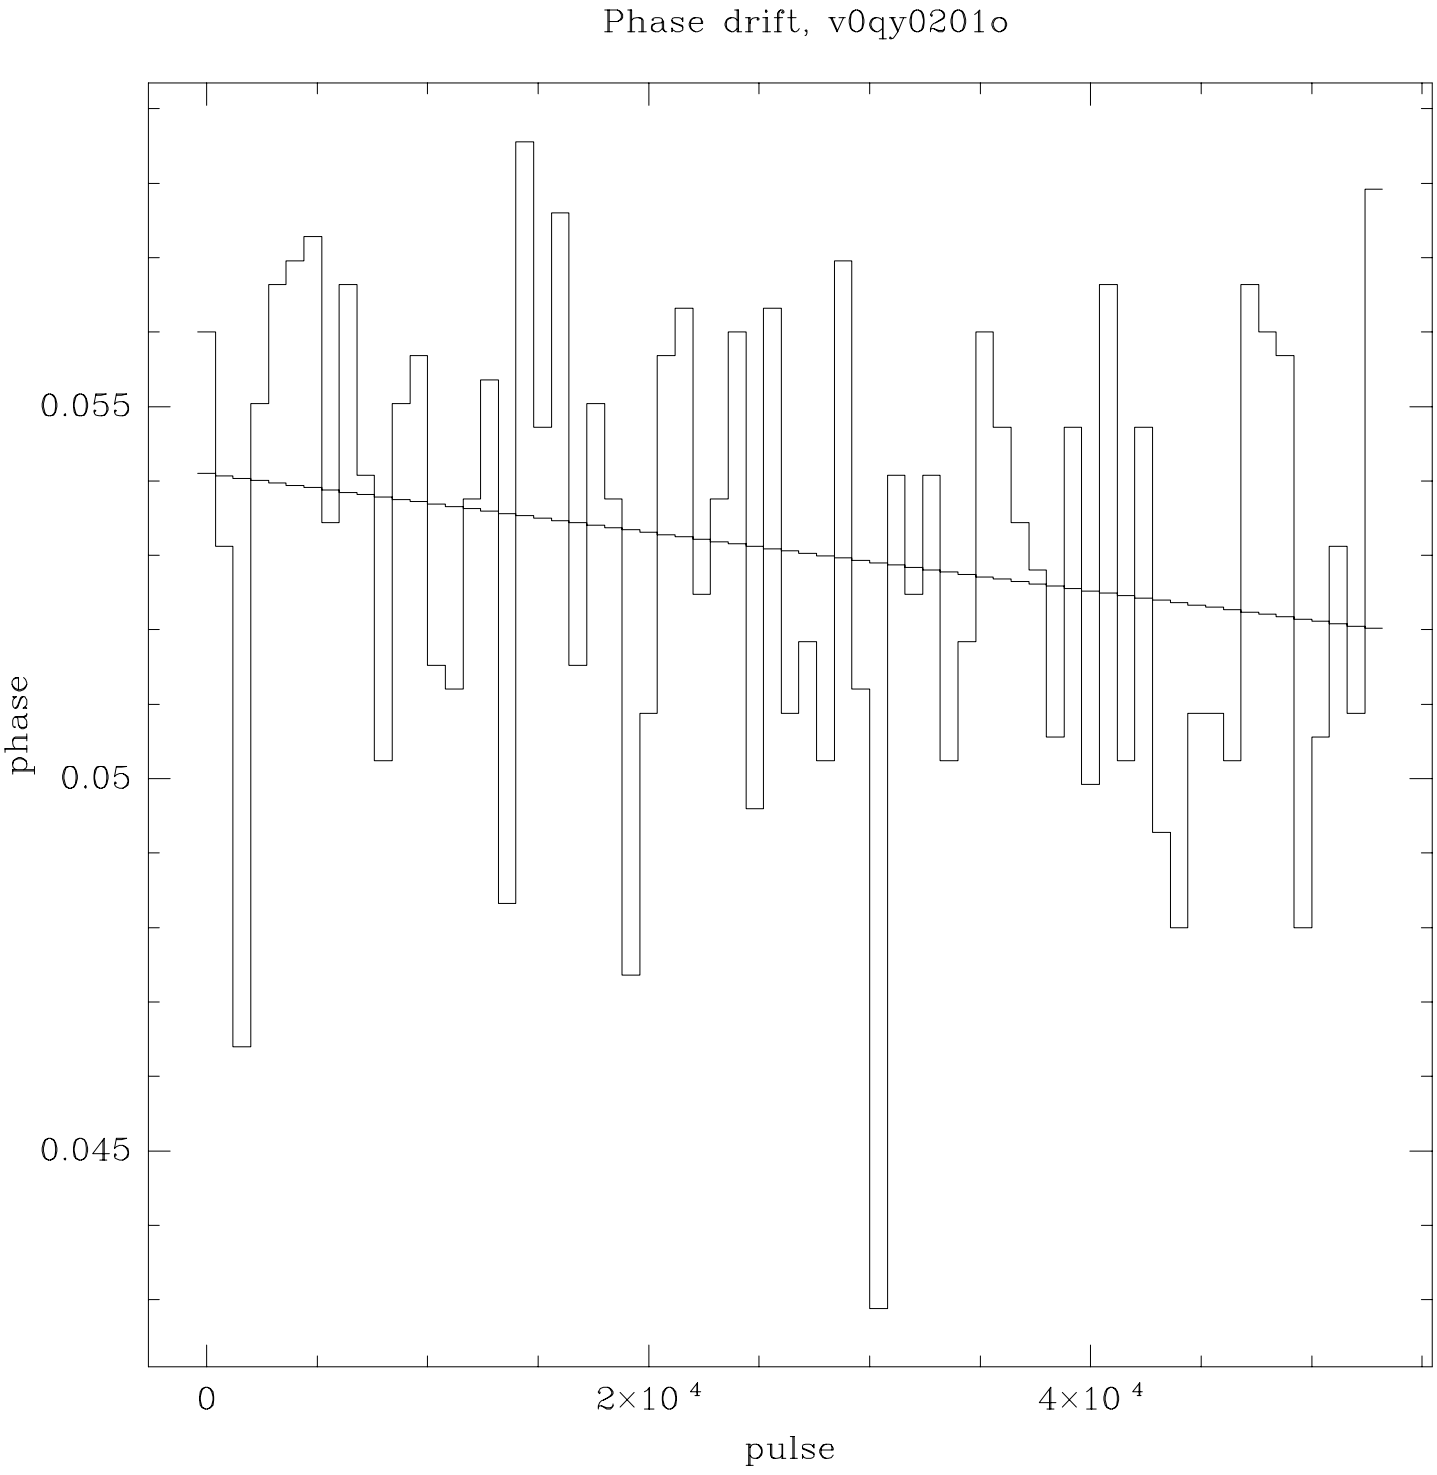

Figure 7:

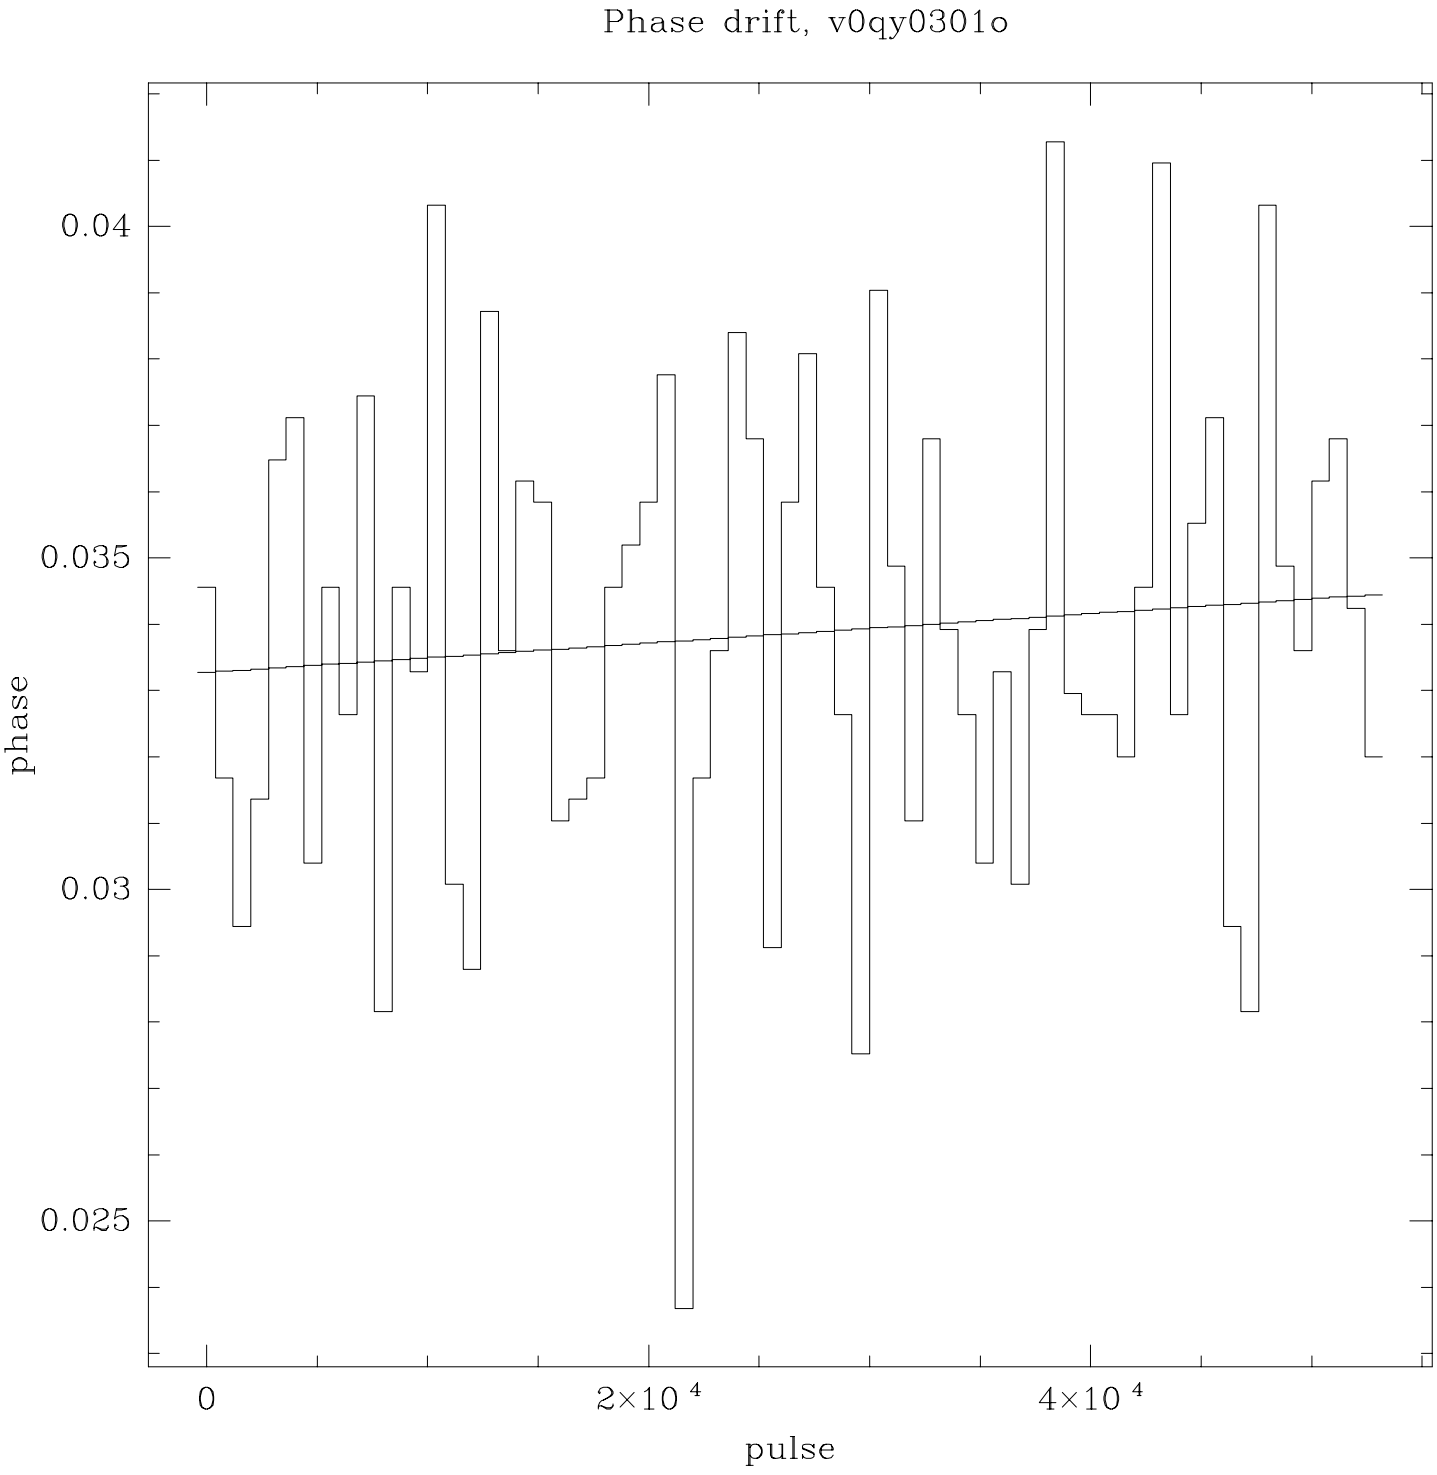
Phase drift, v0qy0301o
0.04
0.035
0.03
0.025
phase
0
2×10 4
4×10 4
pulse

Figure 8:

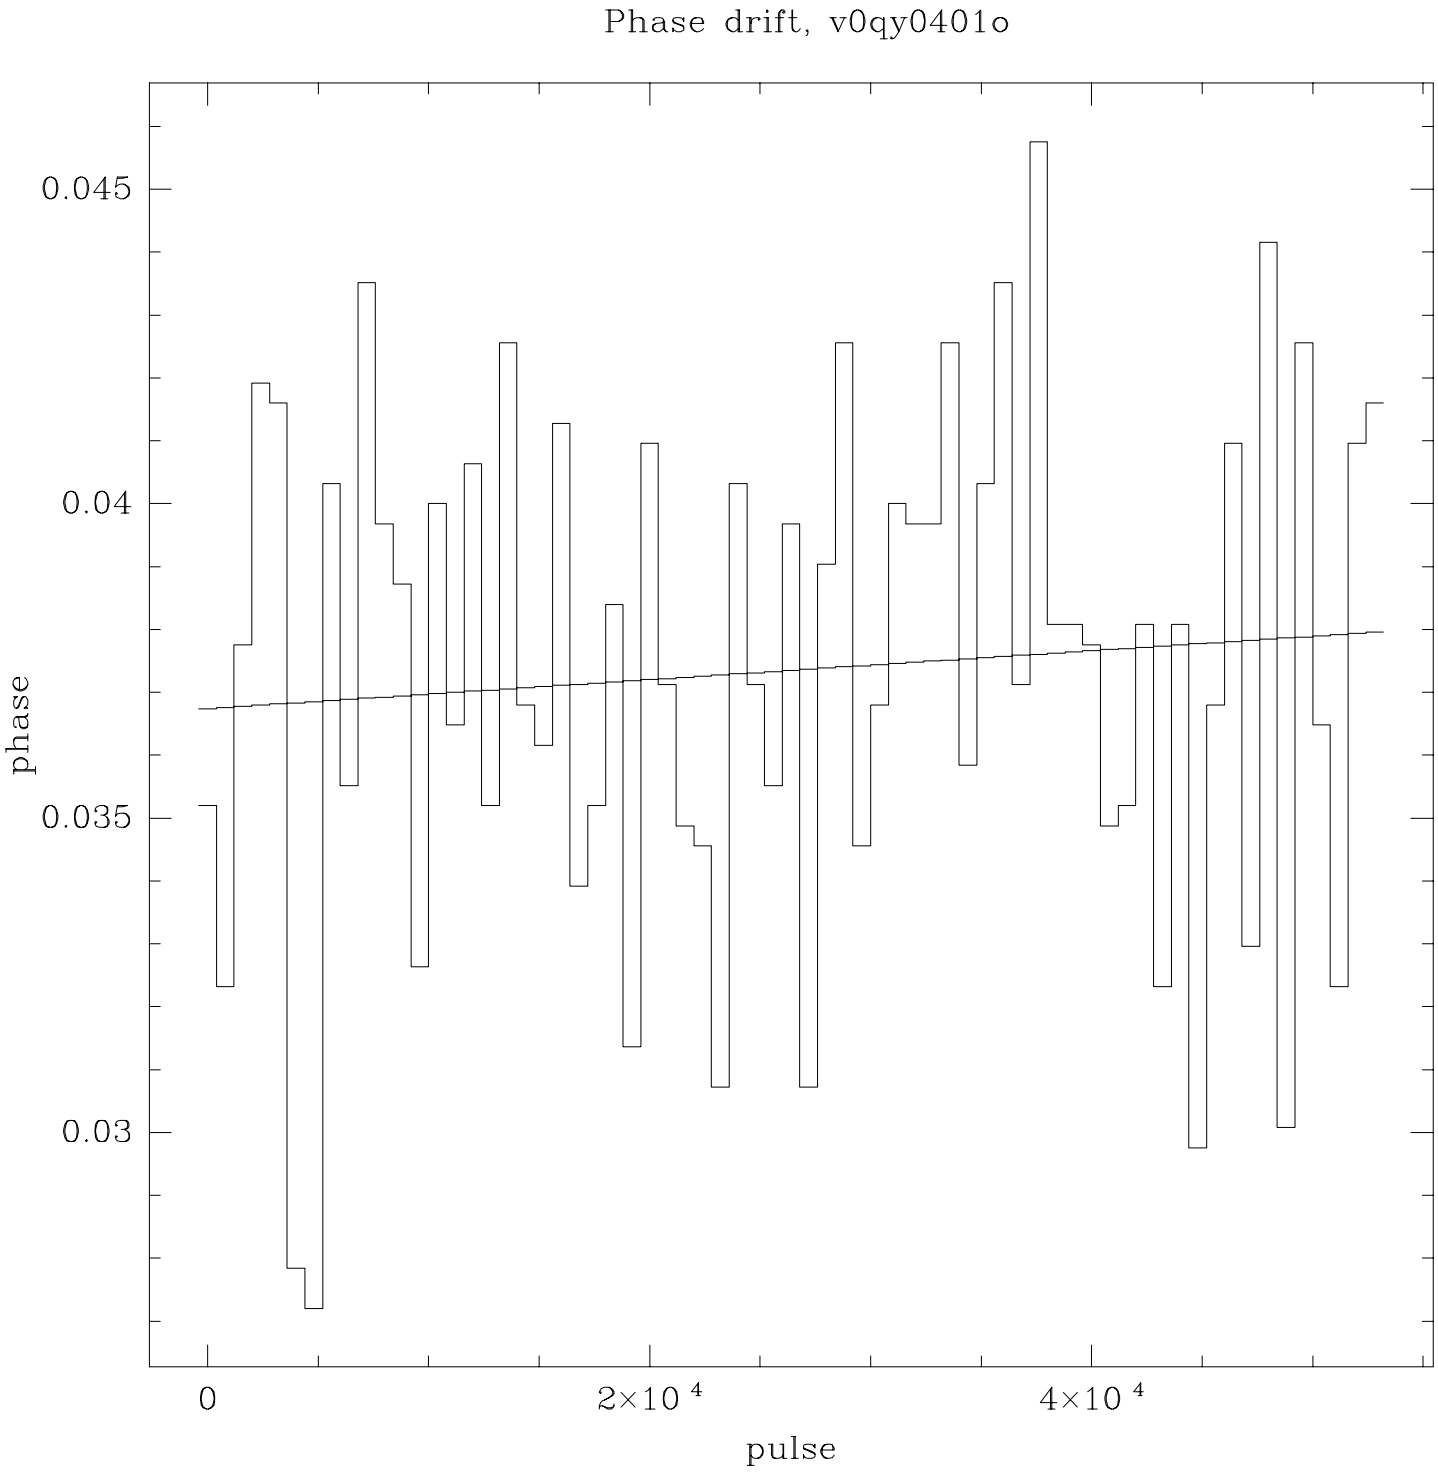
Phase drift, v0qy0401o
0.045
0.04
0.035
0.03
phase
0
2×10 4
4×10 4
pulse

Figure 9:

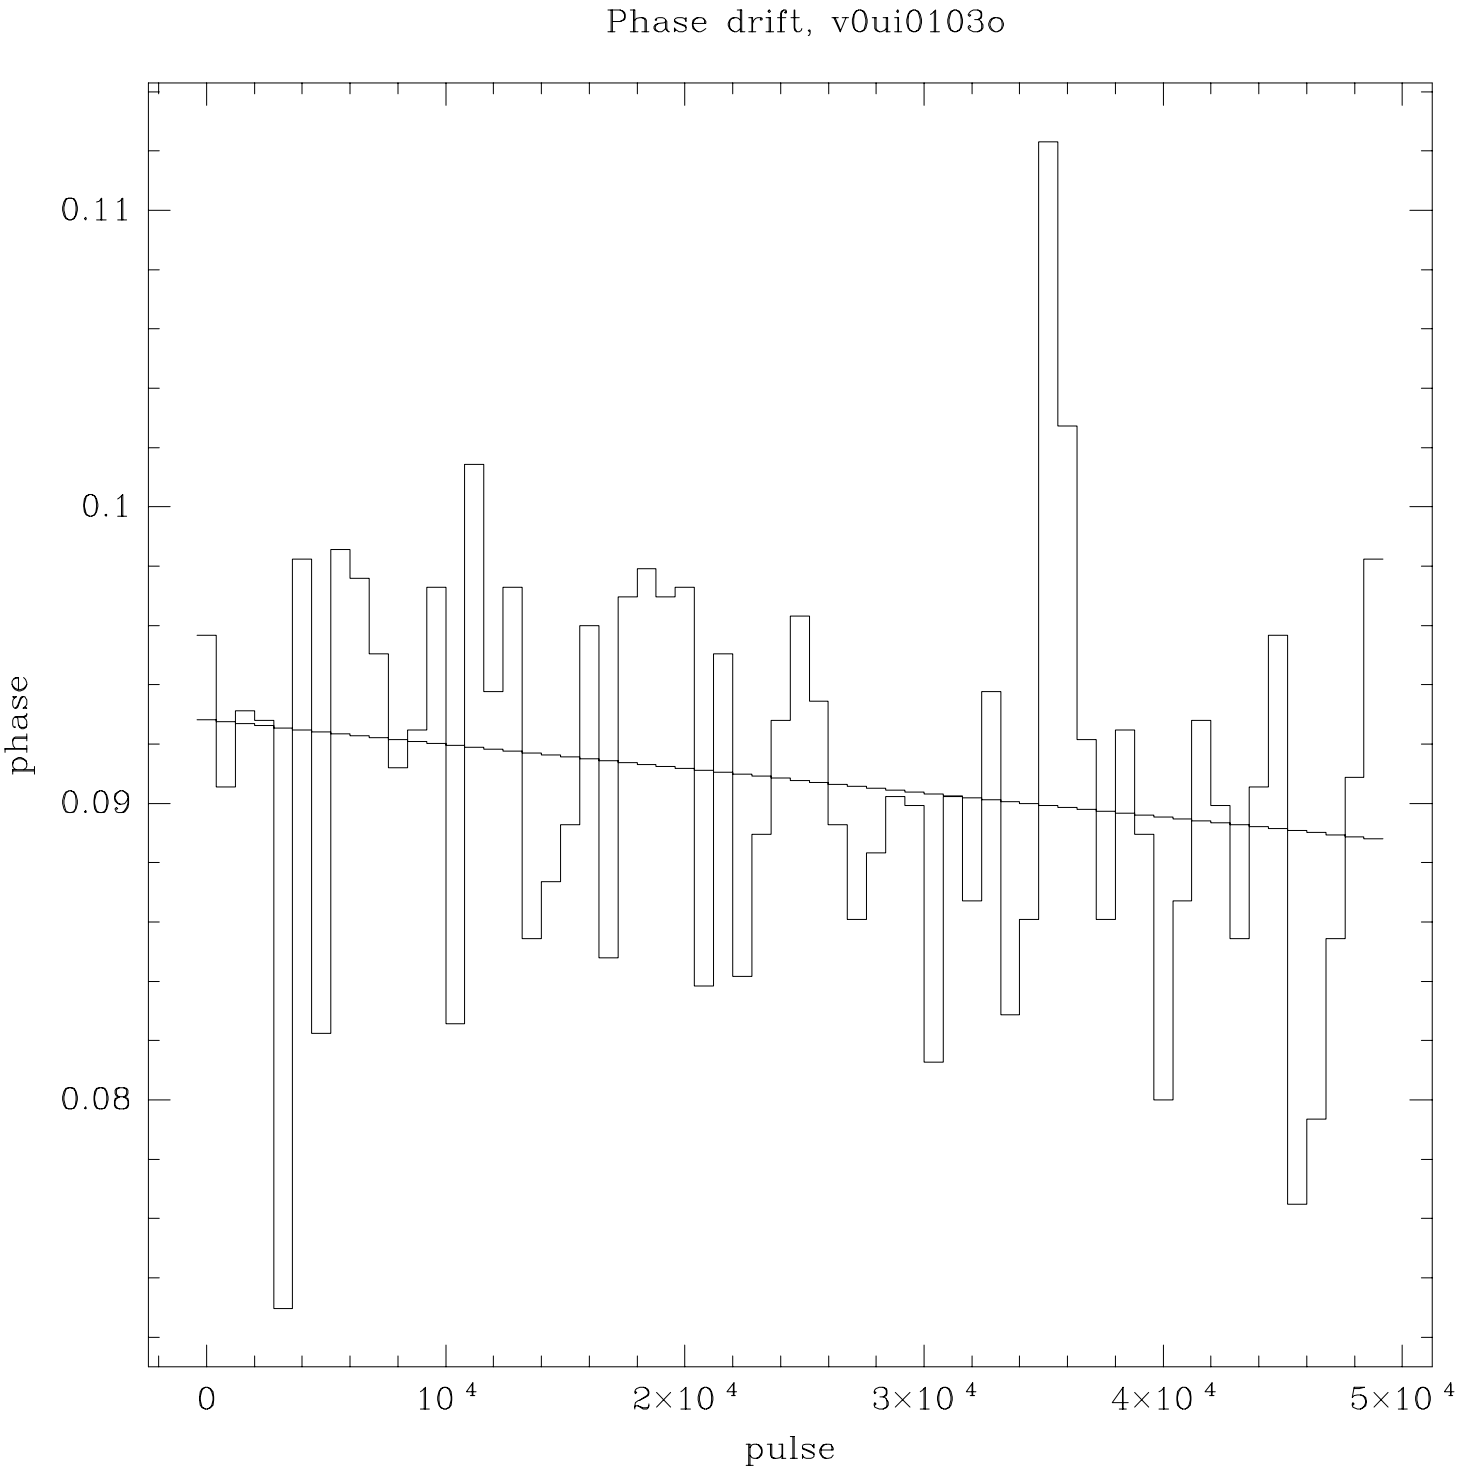
Phase drift, v0ui0103o
0.11
0.1
0.09
0.08
phase
0
10 4
2×10 4
3×10 4
4×10 4
5×10 4
pulse

Figure 10:

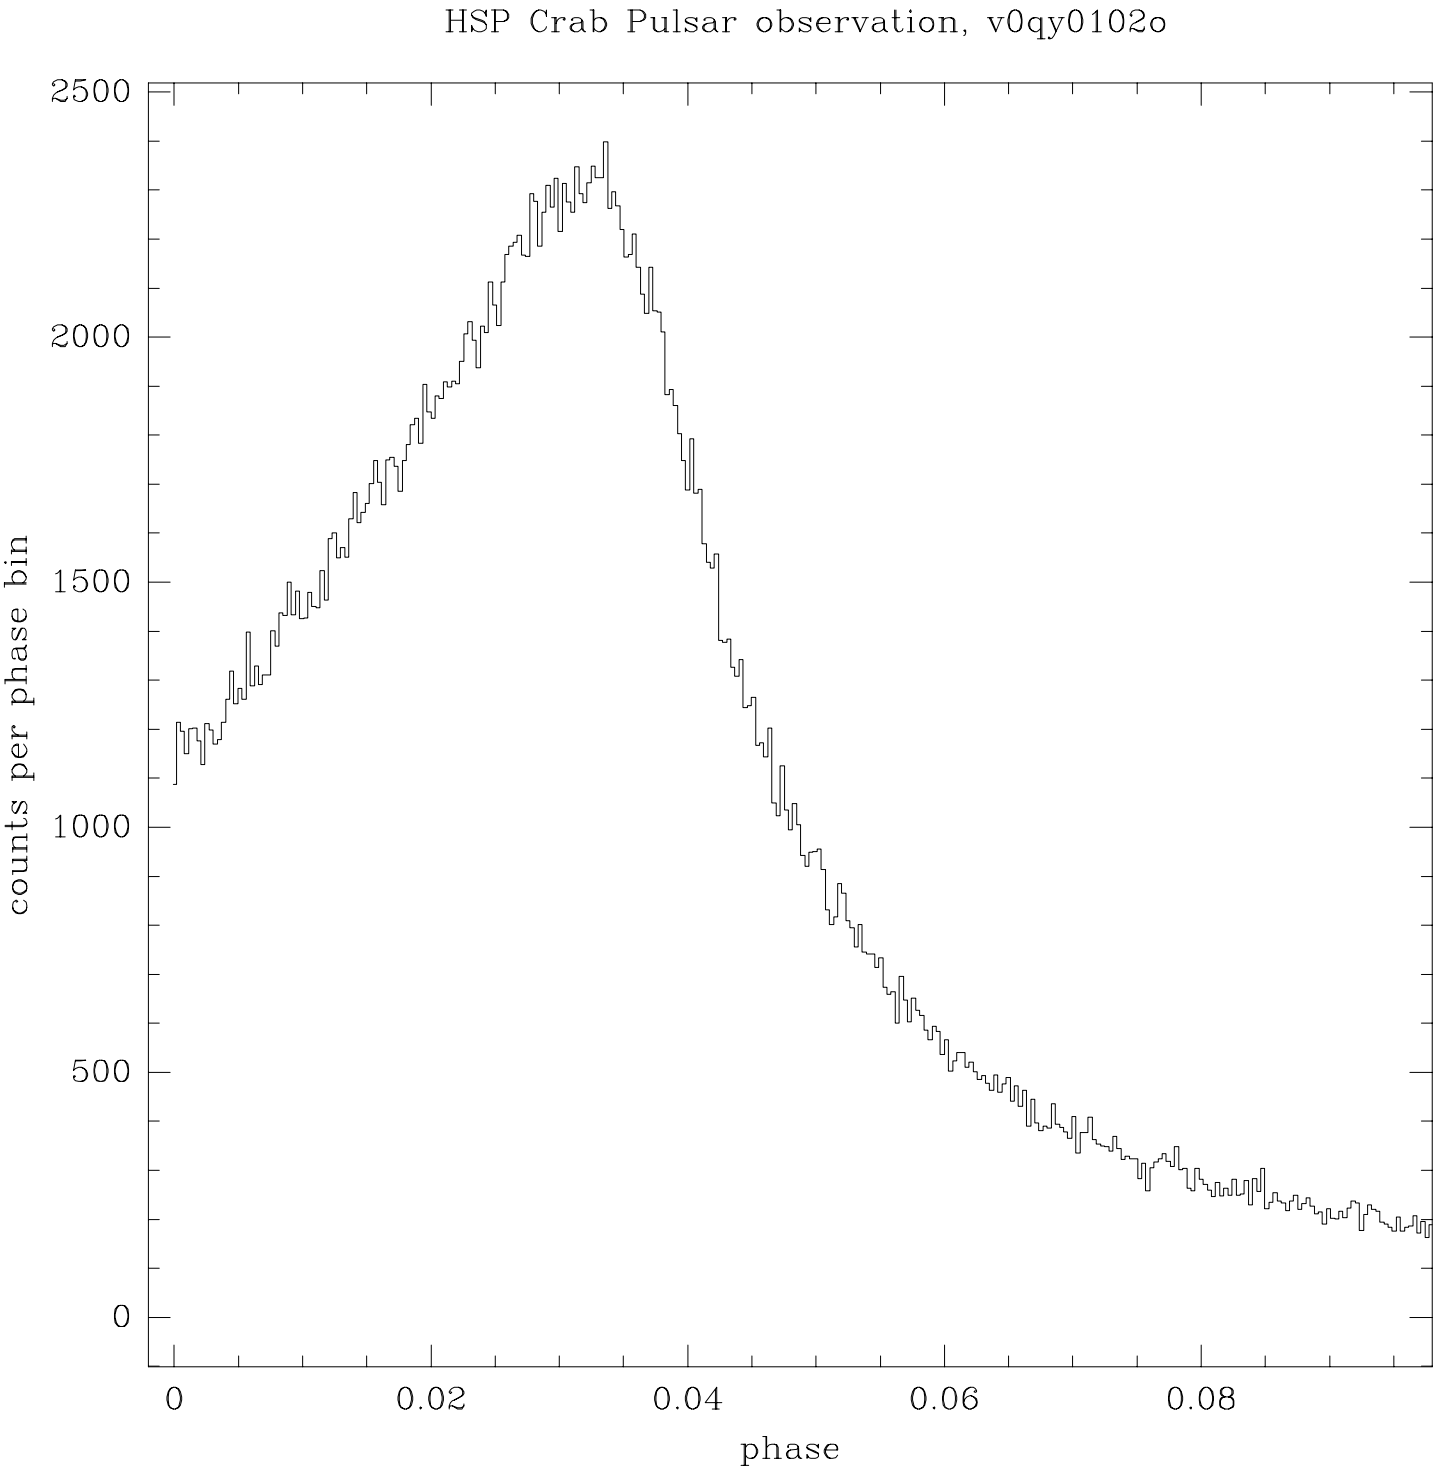
HSP Crab Pulsar observation, v0qy0102o
counts per phase bin
phase
2500
2000
1500
1000
500
0
0
0.02
0.04
0.06
0.08

Figure 11:

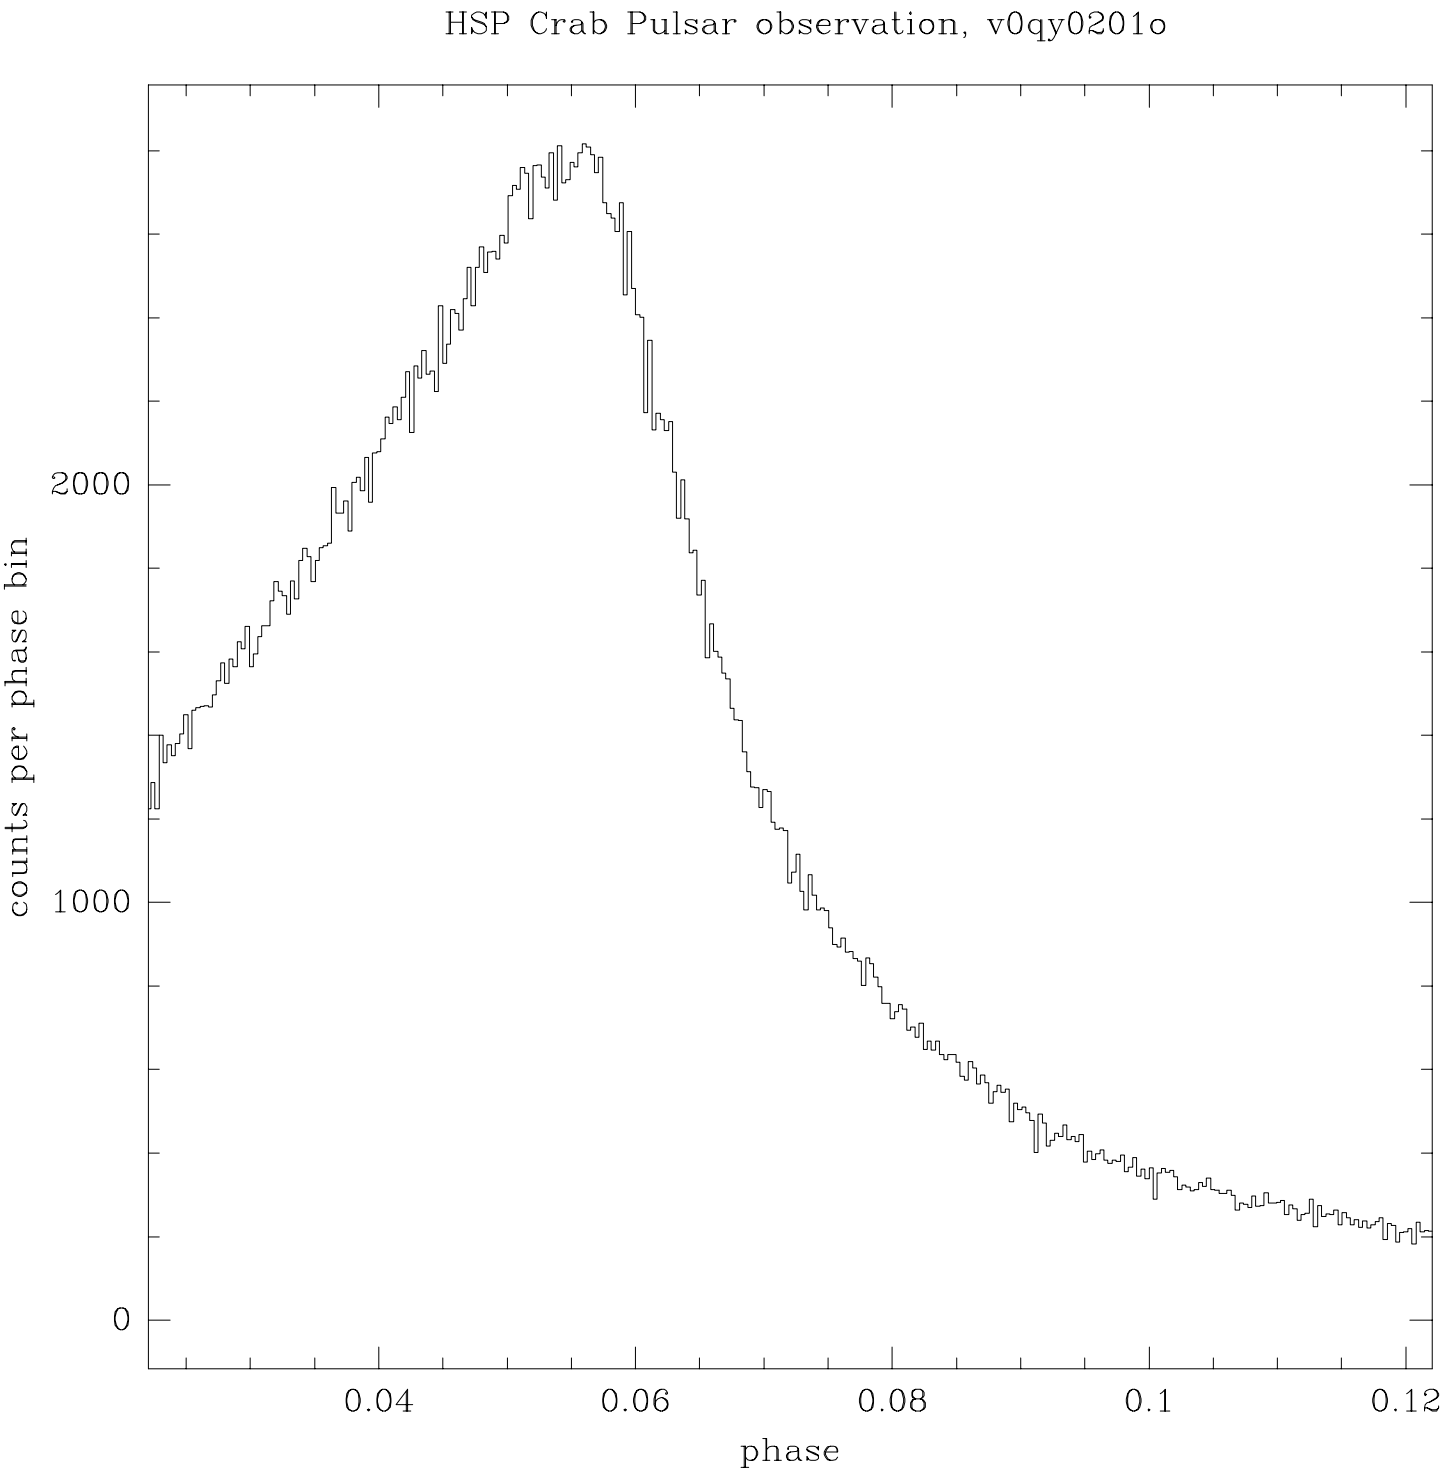
HSP Crab Pulsar observation, v0qy0201o
counts per phase bin
2000
1000
0
0.04
0.06
0.08
0.1
0.12
phase

Figure 12:

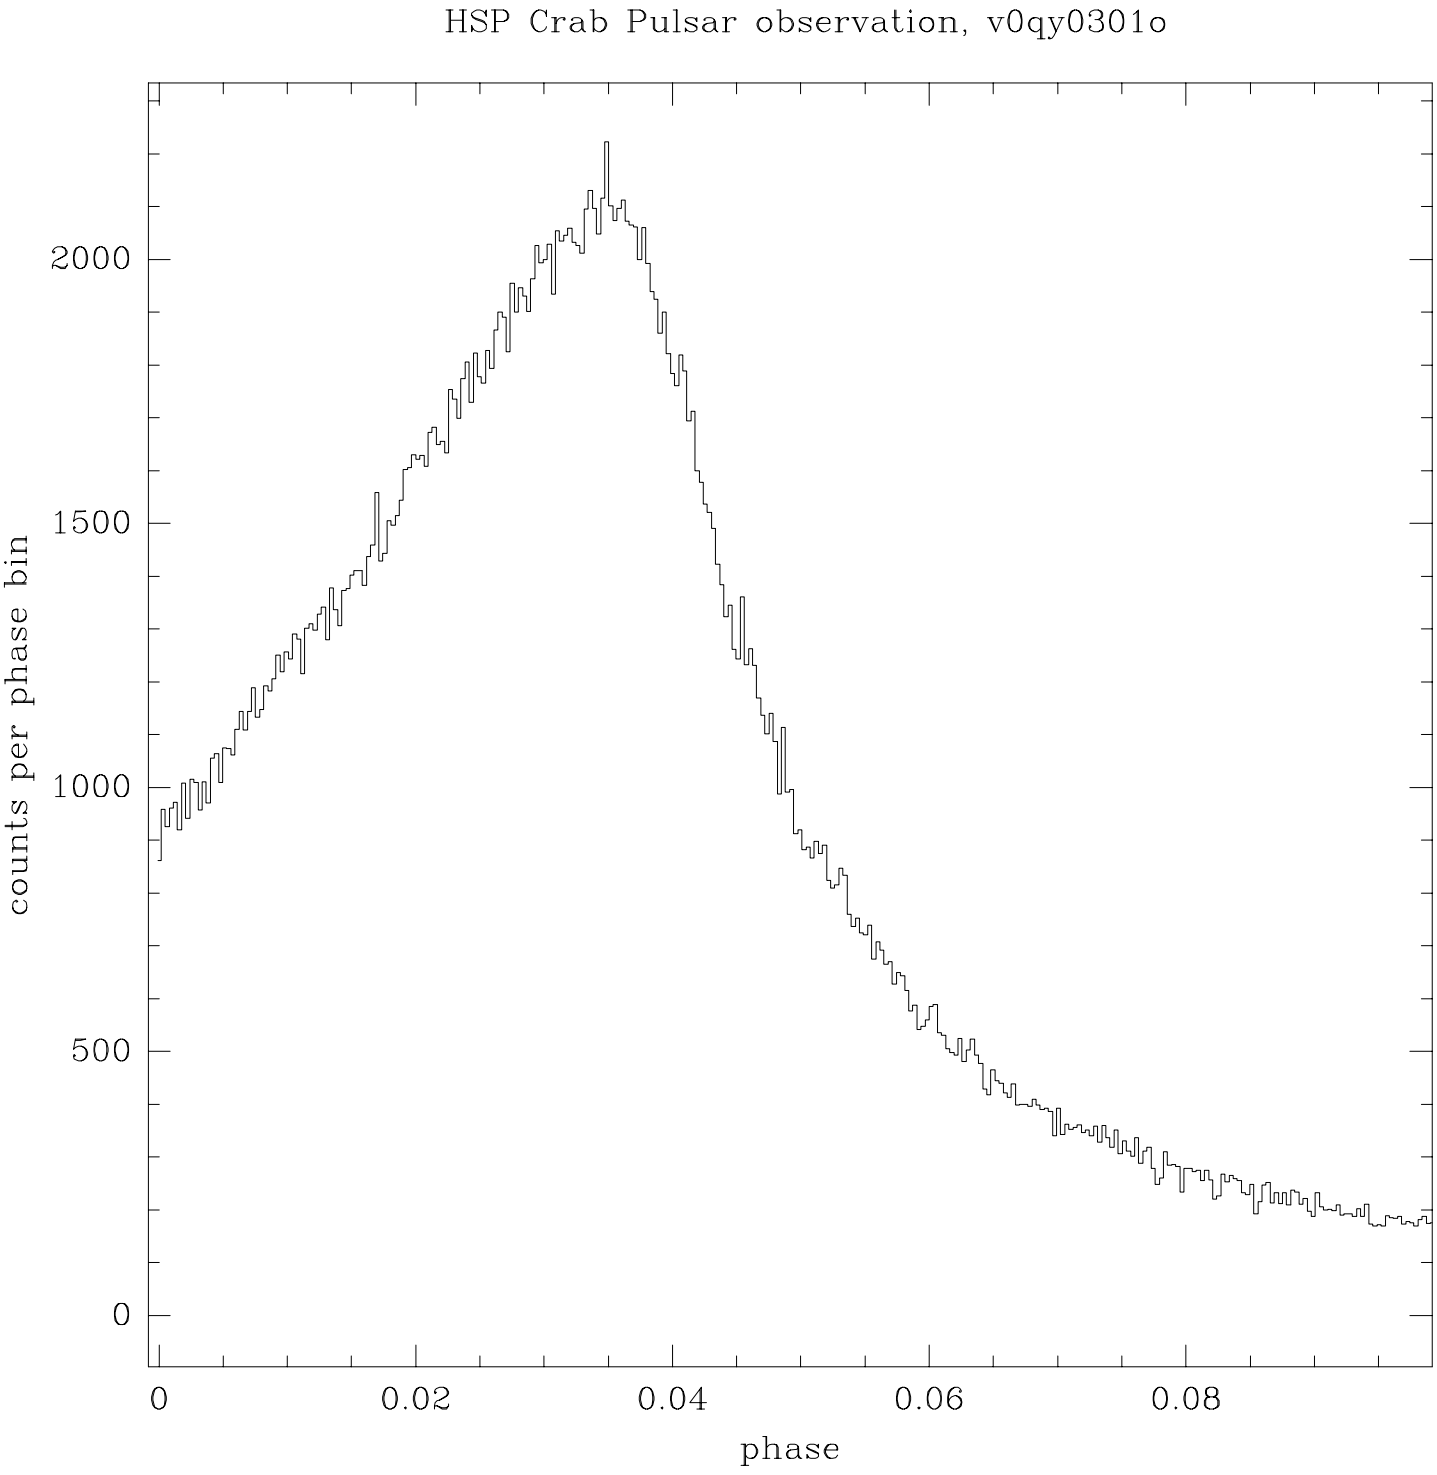
HSP Crab Pulsar observation, v0qy0301o
counts per phase bin
0
500
1000
1500
2000
0
0.02
0.04
0.06
0.08
phase

Figure 13:

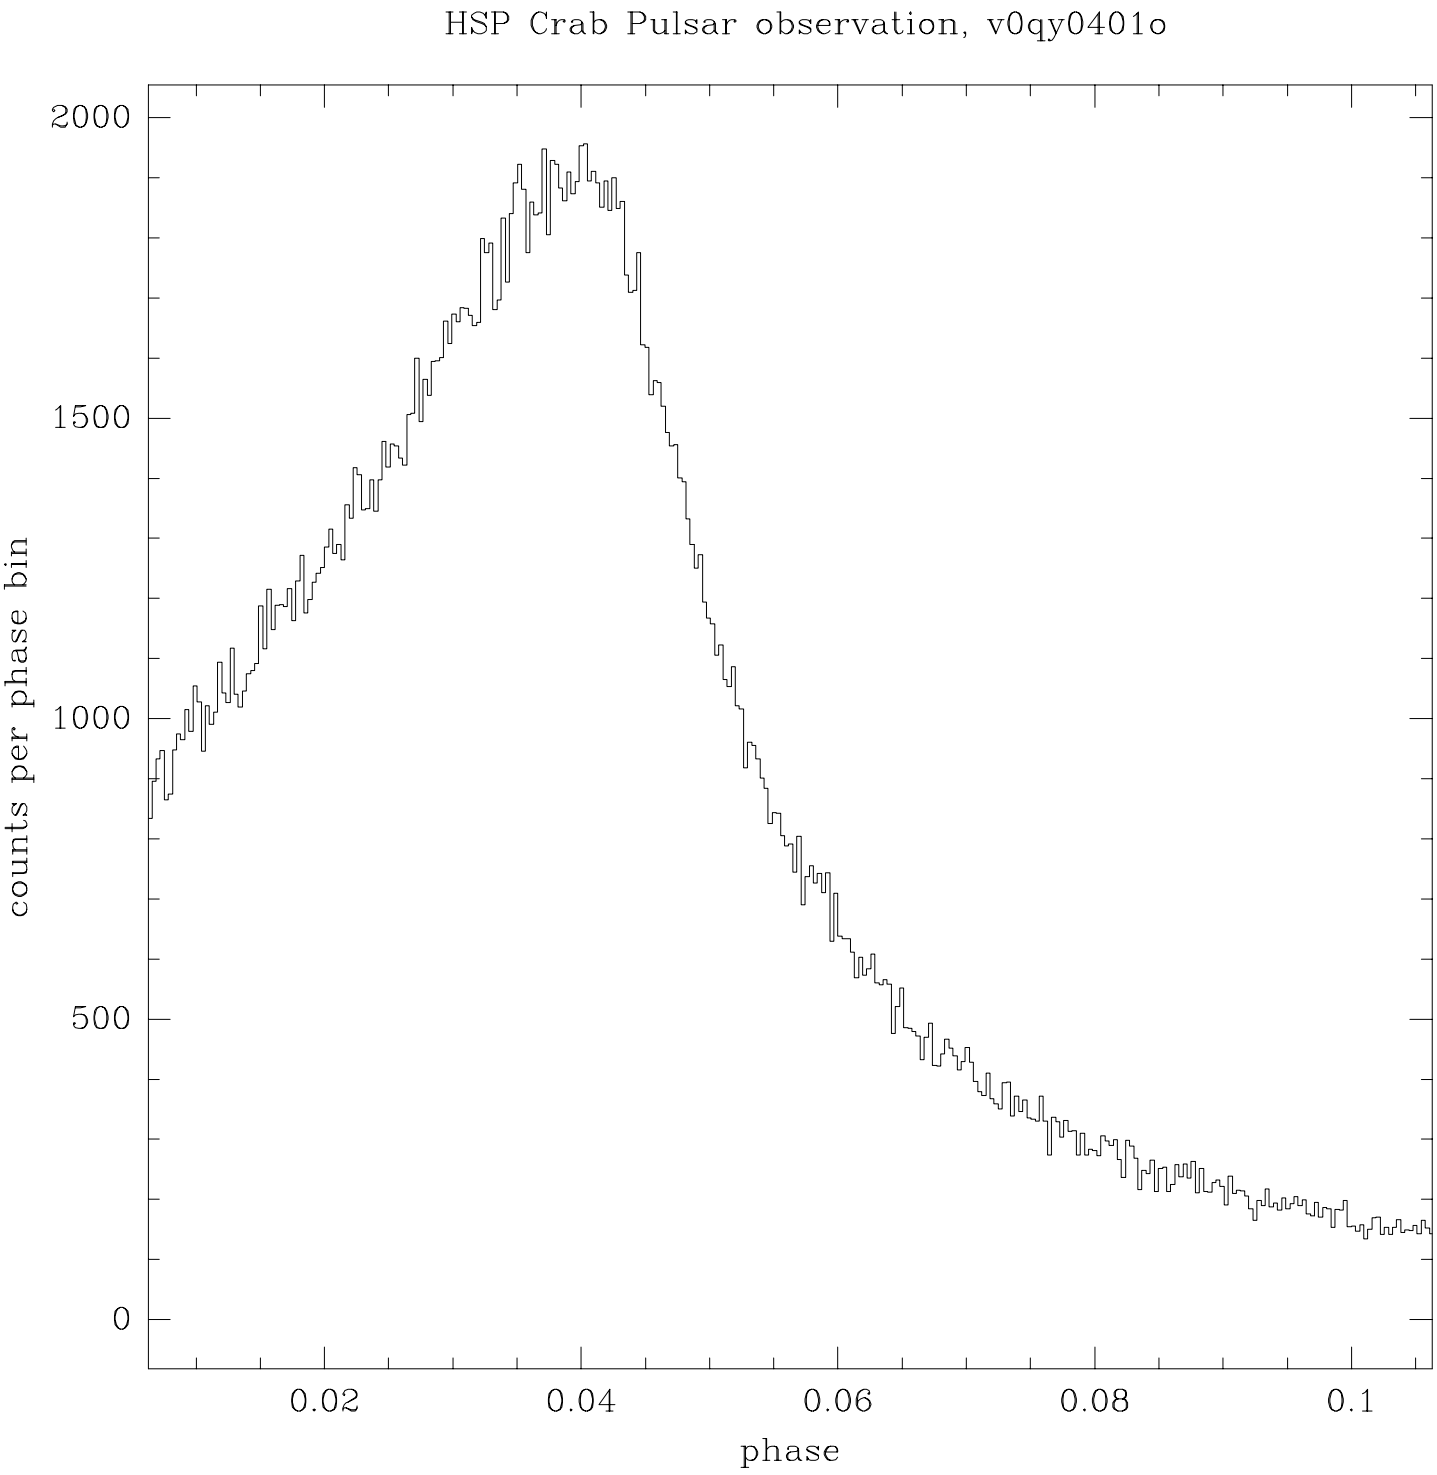
HSP Crab Pulsar observation, v0qy0401o
2000
1500
1000
500
0
counts per phase bin
0.02
0.04
0.06
0.08
0.1
phase

Figure 14:

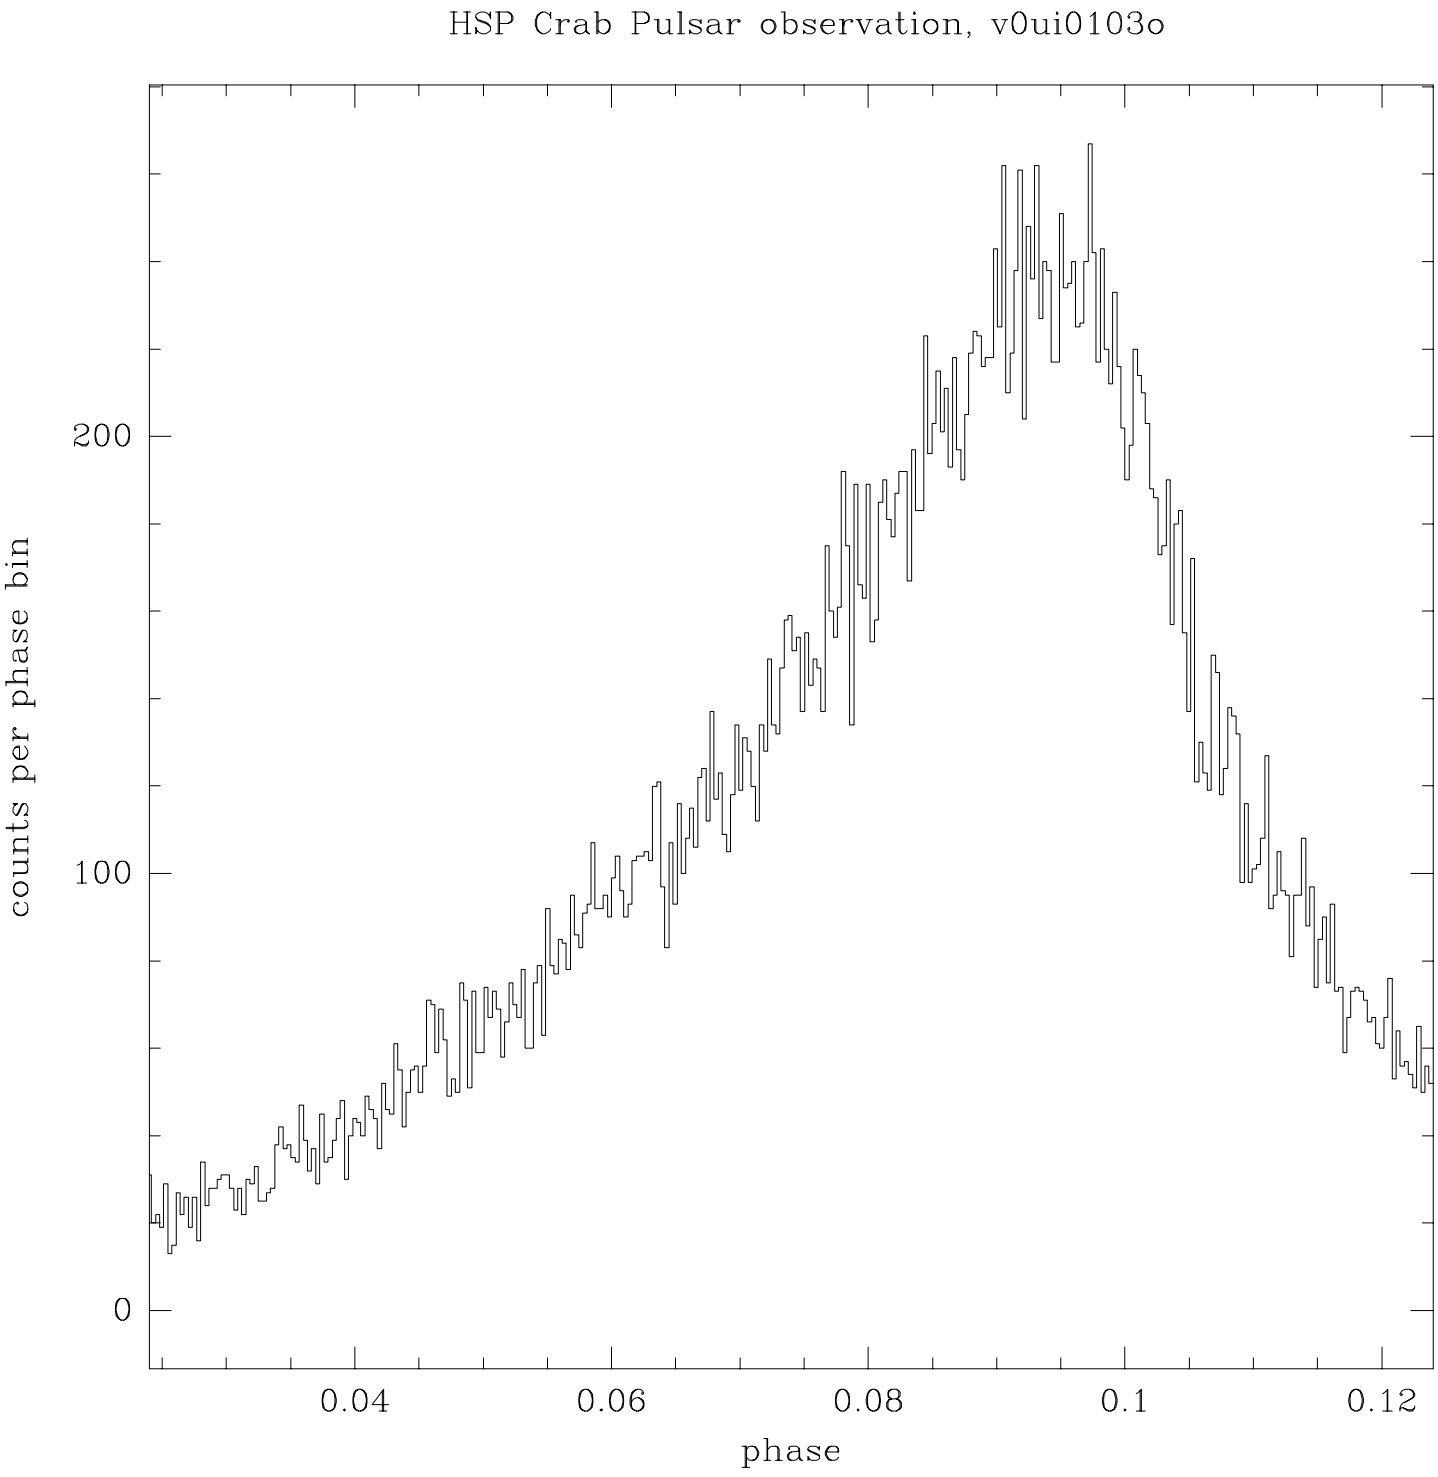
HSP Crab Pulsar observation, v0ui0103o
counts per phase bin
0
100
200
0.04
0.06
0.08
0.1
0.12
phase

Figure 15: