



Publication Year	2021
Acceptance in OA	2022-05-09T14:27:02Z
Title	Timing Calibration of the NuSTAR X-Ray Telescope
Authors	BACHETTI, Matteo, Markwardt, Craig B., Grefenstette, Brian W., Gotthelf, Eric V., Kuiper, Lucien, Barret, Didier, Cook, W. Rick, Davis, Andrew, Fürst, Felix, Forster, Karl, Harrison, Fiona A., Madsen, Kristin K., Miyasaka, Hiromasa, Roberts, Bryce, Tomsick, John A., Walton, Dominic J.
Publisher's version (DOI)	10.3847/1538-4357/abd1d6
Handle	http://hdl.handle.net/20.500.12386/32089
Journal	THE ASTROPHYSICAL JOURNAL
Volume	908



Timing Calibration of the NuSTAR X-Ray Telescope

Matteo Bachetti^{1,2,10}, Craig B. Markwardt³, Brian W. Grefenstette², Eric V. Gotthelf⁴, Lucien Kuiper⁵,
Didier Barret⁶, W. Rick Cook², Andrew Davis², Felix Fürst⁷, Karl Forster², Fiona A. Harrison², Kristin K. Madsen^{2,3},

Hiromasa Miyasaka², Bryce Roberts⁸, John A. Tomsick⁸, and Dominic J. Walton⁹

¹ INAF-Osservatorio Astronomico di Cagliari, via della Scienza 5, I-09047 Selargius, Italy; matteo.bachetti@inaf.it

² Space Radiation Laboratory, Caltech, 1200 E California Blvd, Pasadena, CA 91125, USA

³ Astrophysics Science Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

⁴ Columbia Astrophysics Laboratory, Columbia University, 550 West 120th Street, New York, NY 10027-6601, USA

⁵ SRON Netherlands Institute for Space Research, Sorbonnelaan 2, NL-3584 CA Utrecht, The Netherlands

⁶ IRAP, Université de Toulouse, CNRS, UPS, CNES, 9, Avenue du Colonel Roche, BP 44346, F-31028 Toulouse Cedex 4, France

⁷ European Space Astronomy Centre (ESA/ESAC), Operations Department, Villanueva de la Cañada (Madrid), Spain

⁸ Space Sciences Laboratory, 7 Gauss Way, University of California, Berkeley, CA 94720-7450, USA

⁹ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

Received 2020 September 16; revised 2020 November 27; accepted 2020 December 7; published 2021 February 23

Abstract

The Nuclear Spectroscopic Telescope Array (NuSTAR) mission is the first focusing X-ray telescope in the hard X-ray (3–79 keV) band. Among the phenomena that can be studied in this energy band, some require high time resolution and stability: rotation-powered and accreting millisecond pulsars, fast variability from black holes and neutron stars, X-ray bursts, and more. Moreover, a good alignment of the timestamps of X-ray photons to UTC is key for multi-instrument studies of fast astrophysical processes. In this paper, we describe the timing calibration of the NuSTAR mission. In particular, we present a method to correct the temperature-dependent frequency response of the on-board temperature-compensated crystal oscillator. Together with measurements of the spacecraft clock offsets obtained during downlinks passes, this allows a precise characterization of the behavior of the oscillator. The calibrated NuSTAR event timestamps for a typical observation are shown to be accurate to a precision of $\sim 65 \mu\text{s}$.

Unified Astronomy Thesaurus concepts: Pulsar timing method (1305); Cross-validation (1909); X-ray detectors (1815); X-ray astronomy (1810); X-ray telescopes (1825); Millisecond pulsars (1062)

1. Introduction

The Nuclear Spectroscopic Telescope ARray (NuSTAR; Harrison et al. 2013) is the first focusing hard (3–79 keV) X-ray mission. Compared to other missions covering the same energy range, NuSTAR provided a 10-fold improvement in angular resolution (58"HPD), while also granting good spectral resolution (0.4 keV at 6 keV), and high effective area ($\sim 1000 \text{ cm}^2$ at 10 keV). NuSTAR was designed with an absolute timing accuracy requirement of 100 ms (Harrison et al. 2013). Typically, the uncertainties in the time-stamping of NuSTAR data are much smaller than this requirement (Madsen et al. 2015): most electronic and propagation delays can be tracked down to $\sim 100 \mu\text{s}$ precision. The largest unmodeled issue in the NuSTAR timing calibration has been, until now, a $\sim 2 \text{ ms}$ drift of the spacecraft clock which is only tracked accurately between ground passes, spaced by a few hours. Gotthelf & Bogdanov (2017) showed that the drift could be tracked by using fast millisecond pulsars with very sharp pulse profiles, and that on short timescales the time measurement was stable enough to show $\sim 15 \mu\text{s}$ -thin pulsar peaks. However, the actual process creating the drift remained unknown, limiting the range of timing capabilities of NuSTAR.

In this paper, we show that this drift is largely due to a temperature-dependent frequency drift of the spacecraft temperature-compensated quartz oscillator (TCXO). Our detailed modeling of this behavior and the clock aging allows the timing calibration of NuSTAR to be improved by almost two orders of magnitude. This opens up NuSTAR to a wide range of new

applications involving rapid variability, like rotation- and accretion-powered millisecond pulsars, X-ray bursts, and kHz quasiperiodic oscillations (van der Klis 2006; Lorimer 2008), or precise synchronization with other satellites for multi-instrument and/or multiwavelength studies.

We describe the NuSTAR time tagging procedure in Section 2. Then, we model the temperature dependence of the TCXO in Section 3. We discuss the improved timing calibration in Section 4.

2. NuSTAR Time Tagging Details

The method used to assign photon arrival times on the spacecraft for NuSTAR detector events is described in detail in Madsen et al. (2015), Bachetti et al. (2015). Photon arrival timestamps are assigned to received photons using the on-board spacecraft oscillator. The spacecraft oscillator itself is calibrated to UT using ground station ranging measurements. Converting photon arrival times to UT requires an accurate model of the uncertainties in the on-board electronics, the stability of the spacecraft reference clock, the propagation delays during the downlink of the data, and the precision of the reference clock in the ground station.

The systematic electronic delays on board are modeled to a few μs ; the propagation delay to the ground station is on the order of 300 ms and is modeled accurately ($< 100 \mu\text{s}$) thanks to frequent tracking of NuSTAR's geographic location¹¹; the Malindi ground station, nominally used to measure NuSTAR's

¹⁰ Fulbright Visiting Scholar.

¹¹ NuSTAR uses the same downlink procedure as Swift, that was validated to the $\sim 100 \mu\text{s}$ level (Cusumano et al. 2012).

clock offsets, is synchronized to UTC via GPS with a precision of a few nanoseconds. This is not true of the less frequently used USN Hawaii and Singapore ground stations, that guarantee only 1 ms precision (see below).

However, the largest source of uncertainty in the NuSTAR clock modeling is the stability of the spacecraft’s oscillator. The mission does not carry on board a GPS-synchronized clock. Event timestamps are referred to the pulse-per second (PPS) signal coming from an MC623X2 TCXO. The nominal frequency (or *divisor*) of this oscillator is ~ 24 MHz, meaning that the oscillator ticks ~ 24 million times before sending the next PPS, a number that can be adjusted remotely through a command to the spacecraft. However, this frequency changes by ~ 1 ppm with the change of the oscillator temperature during an orbit. This produces an accumulated delay between the recorded time and the reference UTC time from the ground station.

Every time NuSTAR connects to a ground station for the downlink of the data, the ground station measures the relative departure of the clock time from UT. If the clock has drifted more than an acceptable limit (chosen to be about 100 ms), the divisor of the spacecraft TCXO is adjusted in order to change the direction of the drift. This is done to satisfy the mission requirement of 100 ms timing precision. The clock offset measurements and divisor changes are used by the NuSTAR Science Operations Center (hereafter SOC) to produce the clock-correction files, which are distributed as part of the NuSTAR CALDB from HEASARC and usable with the FTOOL `barycorr` needed to convert spacecraft UT time to the solar system barycenter.

However, the temperature changes rapidly during each of the ~ 97 minutes of a NuSTAR orbit, which is found to effect the TCXO frequency with a similar ~ 97 m modulation. Because of the relatively infrequent ground contacts, the time offset measurements do not accurately reflect the instantaneous arrival times of the events as registered by the spacecraft. As a consequence, the clock correction file produced at the NuSTAR SOC is only able to model the drift to the ~ 2 ms level.

In the following sections, we characterize the frequency-temperature relation in detail, allowing a much better approximation of the clock drift.

3. Temperature and Aging Characterization of the TCXO

3.1. Basic Modeling Framework

The observable we want to measure here is the real TCXO frequency f_0 , which is the reference for time tagging.

The *clock divisor* $D(t)$ is a commanded value, and can be changed as needed by the ground station. The spacecraft 1-PPS tick count rolls over when D clock cycles have occurred. The “spacecraft time” can be thus obtained:

$$\mathcal{T}(t) = \int_0^t \frac{f_0(t')}{D(t')} dt' \quad (1)$$

Periodic measurements from the ground station return the difference between the spacecraft time \mathcal{T}_i and UT time t_i :

$$d(t_i) = \mathcal{T}_i - t_i + \epsilon \quad (2)$$

where ϵ is an error term that we need to estimate. The mission operations center (MOC) removes known biases and delays,

which we exclude from the formula above (but whose uncertainty we include in the error term).

Let us for a moment assume that we know the function $d(t)$ continuously and that we can ignore the error term. Using Equations (1) and (2), we get an estimate of the instantaneous spacecraft clock frequency

$$f_0(t) = \dot{\mathcal{T}}(t)D(t) = [\dot{d}(t) + 1]D(t) \quad (3)$$

where the dotted variables indicate time derivatives. Considering that $\dot{d}(t)$ is measured in discrete intervals, the Equation (3) can be approximated as

$$\bar{f}_0(t) \approx \left[\frac{\Delta d(t)}{\Delta t} + 1 \right] D(t) \quad (4)$$

$$\approx \left[\frac{d(t_2) - d(t_1)}{t_2 - t_1} + 1 \right] D(t) \quad (5)$$

where $\bar{f}_0(t)$ in this case is the *average* frequency between subsequent clock offset measurements $d(t_2)$ and $d(t_1)$. Therefore, by measuring the clock offsets and using the known commanded divisor, we can estimate the frequency change of the clock over a given interval. The right-hand side of Equation (4) is only based on measured and commanded quantities. The left-hand side can be regressed against other quantities averaged over the same time intervals.

For convenience, from now on we will make the calculations not directly using f_0 , but rather its deviation from the nominal TCXO frequency of 24 MHz

$$\mathcal{F}(t) = \frac{\bar{f}_0(t) - 24 \cdot 10^6}{24 \cdot 10^6} \quad (6)$$

Now let us suppose that the TCXO frequency is a function of temperature (an imperfect temperature compensation) and time (clock aging, which makes the quartz crystals less “elastic”; see, e.g., Vig & Meeker 1991). We seek a relation of the form

$$\mathcal{F}(t, T) = \mathcal{A}(T) + \mathcal{B}(t) + \mathcal{C} \quad (7)$$

where \mathcal{A} is only a function of temperature, \mathcal{B} an aging law, unrelated to temperature, and \mathcal{C} is a constant.

We model \mathcal{B} with a function of the kind:

$$\mathcal{B}(t) = b_0 \ln b_1(t - t_0 + 1) + b_2 \ln b_3(t - t_0 + 1) \quad (8)$$

This particular form is motivated by the typical aging behaviors of crystal oscillators, with logarithmic changes of the frequency due to *absorption* of contaminants and *desorption* of crystal particles (Vig & Meeker 1991; Landsberg 1955). In reality, we do not know if these exact processes are driving the aging curve in Figure 1, or if there are other processes involved. But this is the simplest physically motivated model that is able to reproduce the full aging curve, as we will see later. We favor this approach with respect to other approaches that are able to fit the data but need many more logarithmic components and are based on even weaker physical motivation (e.g., Su 1996).

For \mathcal{A} , we use a quadratic function of the temperature:

$$\mathcal{A}(T) = a_0(T - T_0) + a_1(T - T_0)^2 \quad (9)$$

where a_i are the fit coefficients.

In practice, we split the constant \mathcal{C} into two parameters c and e that we used to fit separately for \mathcal{A} and \mathcal{B} .




Figure 1. Correction (in *parts per billion*, ppb) for the long-term degradation of the spacecraft’s crystal oscillator clock over time. (Left) a quadratic fit to the clock frequency in the f_0 - T plane after correcting for aging effects as described in the text. The blue line shows the quadratic fit to Equation (9), while the green line shows the linear component of the model to guide the eye; (right) the aging of the crystal oscillator corrected for temperature effects; bottom panels show the residuals from the above models, whose scatter is less than 10 parts per billion.

To calculate the expected clock offsets (Equation (2)) we iterate the best-fit values for parameters a_i and b_i , above, eliminating the effects of the aging law from the temperature law and vice versa.

The temperature of the TCXO is measured every ~ 10 s and recorded as part of the mission housekeeping information, so there are typically thousands of temperature measurements between ground segment passes. Given a set of temperature measurements T_k for times t_k , the numerical integral becomes

$$\Delta d(t_i) = \sum_{k=0}^i \left\{ \left[1 + \mathcal{F}(t_k, T_k) \right] \frac{24 \cdot 10^6}{D_k} - 1 \right\} (t_{k+1} - t_k) \quad (10)$$

where D_k is the value of the divisor at time t_k . We can then model these offsets with a continuous function (e.g., a spline fit or a more robust interpolation) and compare them with the measured offsets.

The comparison between the measured offsets and the ones obtained in Equation (10) shows two very distinct problems with the clock offset measurements from the ground stations (see Figures 2 and 3):

1. Only the Malindi station has a symmetric scatter around the best solution, while the Singapore and USN Hawaii ones are systematically overestimating the offset. This is

a known effect, due to the fact that the latter two stations truncate their reference clock to 1 ms precision.

2. A large number of measurements, regardless of the ground station, underestimate the offset (the effect is so large that these bad measurements are often easy to flag just by comparing them with nearby measurements). This is most likely due to processing delays within the spacecraft computer during the measurement, probably due to high CPU utilization. Mission operations staff make a best effort to perform clock calibration during periods of expected low CPU utilization, but this goal is not always achieved.

These outlier clock measurements of both kinds can be easily flagged. After these measurements have been removed, we can go back to the beginning and fit Equations (9) and (8) to the cleaned data.

In summary, the fit procedure is iterative, as follows:

1. Fit the offset data to Equation (9) with an initial value of the offset constant c ;
2. Fit Equation (8) plus an offset e a first time to the data after subtracting the results of 1;
3. Flag all points at more than 0.02 parts-per-million (ppm) from the solution as outliers;
4. Fit only the linear term of Equation (9) to the data after subtracting the aging solution;

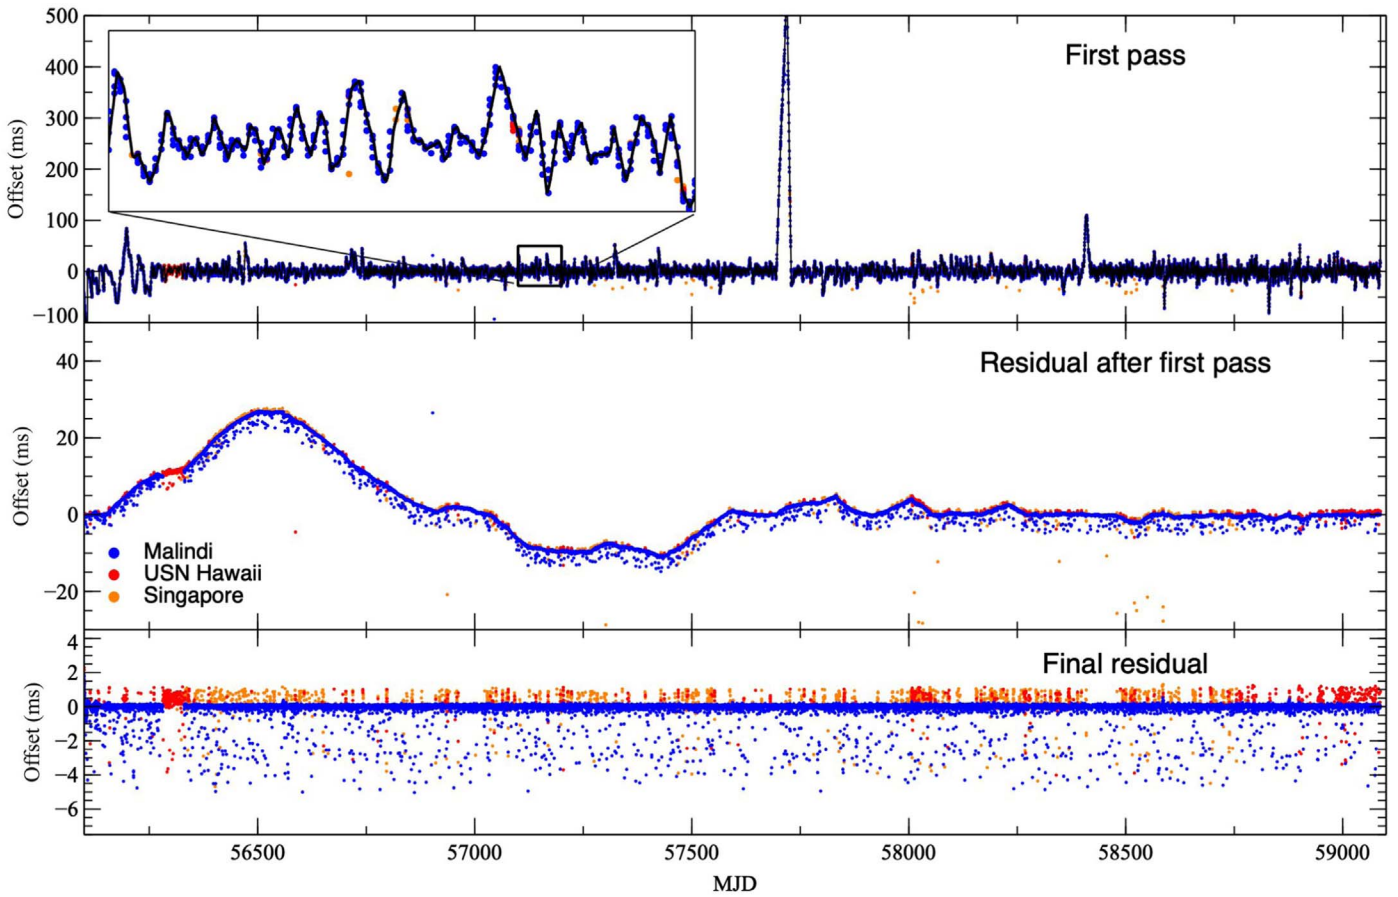


Figure 2. Time history of the drift of the spacecraft clock. Offsets are color-coded to indicate different ground stations, as shown in the legend. (Top) spacecraft clock offsets measured during ground passes, compared to the offset calculated with Equation (10) (in black; the inset shows MJDs 57100–57200); (middle) residuals after correcting for the temperature model and adjusting for major trends. Note the large number of outlier measurements below a smooth long-term trend; (bottom) residuals after eliminating the long-term trend with a robust spline fit.




Figure 3. Histogram showing the scatter of detrended residuals for the three ground stations. Note the much better performance of Malindi with respect to the other ground stations, with a scatter of $\sigma \approx 65 \mu\text{s}$ and an additional “tail” of outlier measurements. The “flat” distribution of Singapore and Hawaii offsets is due to the way these ground stations provide clock offset measurements, truncated to 1 ms precision. Since values are truncated and not rounded, the offset is systematically shifted in one direction.

5. Fit only the quadratic term of Equation (9) to the data after subtracting the aging solution and the linear term. This allows one to constrain precisely T_0 thanks to the symmetry of the quadratic term;
6. Fit the full model A to the data one last time after subtracting the aging solution and fixing T_0 ;
7. Fit Equation (8) one last time on the temperature-detrended data;
8. Calculate the clock offset history using Equation (10); flag outlier clock offset measurements;
9. Go back to point 1 and use the cleaned clock offset history to get more meaningful estimates of the local oscillator frequency. This time, calculate confidence intervals for the free parameters in Equations (9) and (8).

The results of the fit are presented in Table 1 and Figures 1–4. We executed the fit with `lmfit` (Newville et al. 2014). The confidence intervals on model A were calculated through a Monte Carlo Markov Chain using the `emcee` library (Foreman-Mackey et al. 2013), with uniform, unbounded priors, and the recommended burn-in and chain-length. The high correlation of the parameters of model B (all but e) did not allow for meaningful confidence intervals with this method.

Table 1
Best-fit Coefficients in the Clock Correction Model (Equations (7)–(9))

Coefficient	Unit	Value
T_0	C	13.440(25)
t_0	MET	77509250 (fixed)
c	ppm	13.91877(8)
a_0	ppm K ⁻¹	-0.07413(7)
a_1	ppm K ⁻²	0.00158(4)
b_0	ppm	0.00829 ^a
b_1	1 yr ⁻¹	100.26 ^a
b_2	ppm	-0.2518 ^a
b_3	1/yr ⁻¹	0.0335 ^a
e	ppm	-0.0202(1)

Note.

^a Uncertainties on the quantity b_i are unreliable, due to the very high correlations between parameters. Therefore, we fixed them to the best fit in order to get a meaningful confidence interval on the offset parameter e , which is important to set the offset to the full model in Equation (7).

3.2. Comparison with Clock Offsets

The results of the fitting procedure described above are shown in Figure 2. The offsets are extremely well modeled throughout the full lifespan of NuSTAR. The difference between the calculated and observed offset is very small and produces an accumulated drift of ~ 200 ms in six years, which accounts for an error on f_0 of a part in a *billion*. This long-term trend can be removed easily and is not shown in the figures. Shorter-term trends remain, a slow drift of $\lesssim 10$ ms yr⁻¹. The results are so precise that we can single out most of the points where the ground stations had a slow response, and we can easily take them out (see also Figure 3). The remaining trends do not show any obvious correlation with temperature or other observables, but they can be interpolated with a spline. We can now track the clock offsets on timescales of seconds instead of hours, and this allows a correction of the event arrival times to an unprecedented precision. The comparison of the performance of the method without using the clock correction file, with the legacy clock correction, and with the new method, are shown in Figure 5.

4. Discussion: Clock Stability and Precision

4.1. Relative Clock Stability

Figure 6 shows the clock stability during an observation. When we compare the width of the pulse profiles with those tabulated in the literature, they are within $20 \mu\text{s}$ of the previously published values.

To get a more quantitative estimate of the residual clock drift after the temperature correction, we use the available data from the pulsars PSR B1821-24A and PSR B1937+21, chosen for their fast rotation and sharp X-ray pulse profiles (Saito et al. 1997; Takahashi et al. 2001). We split the data set into 97 min long segments (approximately an orbit), and fold the photons in each segment using the X-ray timing solution from the NICER mission (Arzoumanian et al. 2014), calculated by Deneva et al. (2019). We calculate the position of the peak in each data segment using the `fftfit` algorithm (Taylor 1992). Finally, we measure the r.m.s. variation of the pulse phase around the median value.¹² We find the standard deviation of the pulse phase during each of the

¹² We expect the median value to be different from zero, due to the uncertainty in the ground station offset measurement.

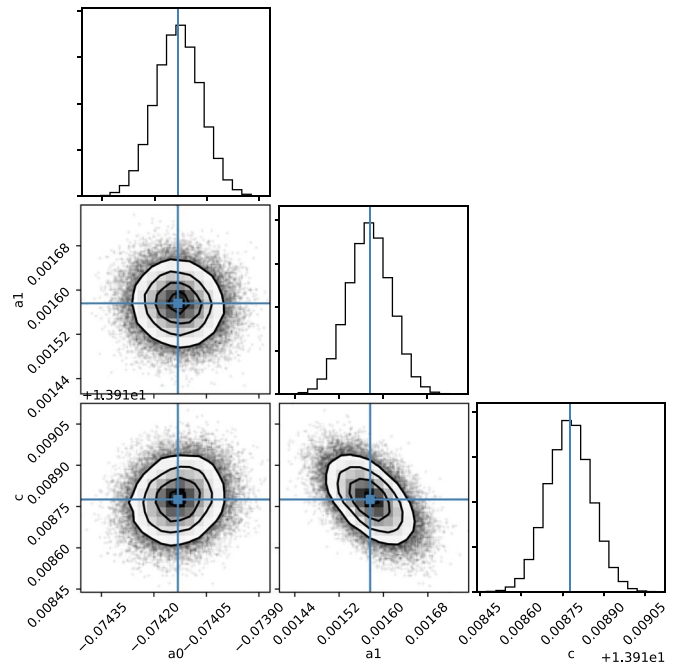


Figure 4. Modeling the imperfect temperature compensation of the spacecraft clock. Confidence intervals for the parameters of Equation (9), plotted using the corner library (Foreman-Mackey 2016).

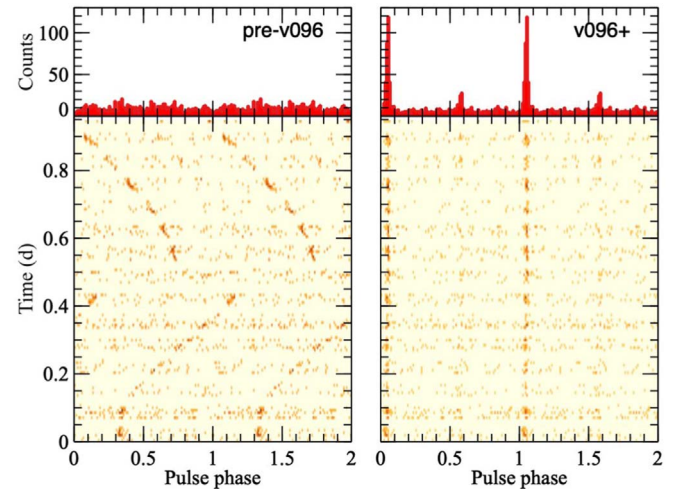


Figure 5. Comparison of the timing performance obtained using the pre-v096 clock correction file (left) and that corrected with the technique described in this paper (right). Shown is the time history of the pulse profile of the rotation-powered millisecond pulsar B1937+21, folded on the ephemeris of Arzoumanian et al. (2018).

long observations of B1821-24A to be between 17.1 and $18.1 \mu\text{s}$, compatible with the width of the main peak of the pulse profile as measured by Deneva et al. (2019). Doing the same exercise with B1937+21, we find a much smaller jitter of $9 \mu\text{s}$. However, we caution that this observation is much shorter, and the number of points is very small. This can produce deceptively small values of the standard deviation.

4.2. Absolute Timing Alignment

As described in Section 4, the clock stability over ~ 1 day of observations is $\sim 10 \mu\text{s}$. We now estimate the *absolute*




Figure 6. Relative stability of the clock, shown through the phaseogram of three observations of PSR B1821-24A and one of PSR B1937+21. To guide the eye, we plot in red a $40 \mu\text{s}$ interval around the peak of the pulsar. The peak of PSR B1821-24A is $\sim 20 \mu\text{s}$ wide, and the additional wandering of the peak in timescales of ~ 1 day is comparable to it. The peak of PSR B1937+21 is thinner, and it stays inside the $\pm 20 \mu\text{s}$ interval during a 1 day long observation. In longer observation, the effect of the uncertainty on ground station offset measurements becomes visible, and the pulse phase shifts by a larger amount, compatible with the $\sim 60 \mu\text{s}$ measured in Figure 3.

timing calibration, i.e., the precision of the NuSTAR clock with respect to a UTC reference. This number is straightforward to measure by looking again at the clock offsets as measured from the ground station passages, and comparing them to the offsets calculated from the model in Section 3. Figure 5 shows how the additional residuals between the Malindi-measured offsets and the temperature model are easily modeled through a spline. Once we do that, it becomes clear that the bulk of Malindi clock offsets are measured to a precision of $\sim 60 \mu\text{s}$. Additionally, it becomes obvious that a large number of clock offset measurements are anomalous, possibly the result of delays in the processing of the data from the spacecraft. These outliers constitute a small fraction of the total measured offset and are easily separated from the “good” measurements. What remains is a quasi-Gaussian distribution of clock offset measurements around zero, with standard deviation $65 \mu\text{s}$, that defines the theoretical long-term reconstructed accuracy of the NuSTAR event timestamps.

However, clock offsets measured by the ground station are not independent of electronic and instrumental delays, and might contain additional unmodeled biases. The ultimate verification of the proper functioning of the NuSTAR clock correction procedures is to compare signals from fast pulsars measured with NuSTAR to those obtained by the NICER mission that has a well-established absolute timing precision of < 300 nanosec. We have selected a sample of fast (< 70 ms) pulsars with sharp features in their pulse profiles for which there exist (quasi) simultaneous NICER (0.2–12 keV) observations. This sample includes three rotation-powered pulsars, one recycled millisecond pulsar, and four accreting millisecond X-ray pulsars, and is described below. In total, we compared 15

simultaneous NICER–NuSTAR observations obtained between 2017 April and 2020 February.

4.2.1. Source Selection

From the soft gamma-ray pulsar catalog provided in Table 2 and Figure 27 of Kuiper & Hermsen (2015), we selected a set of three rotation-powered pulsars satisfying our requirements: PSR B0531+21 (Crab pulsar; 33.5 ms), PSR J0205+6449 (65.7 ms) and PSR J2022+3842 (48.6 ms). The latter two pulsars are very weak at radio frequencies and are currently being timed by NICER as part of an ongoing monitoring program (see e.g., Section 4.1 of Kuiper & Hermsen 2009, for the method using time-of-arrival (TOA) measurements). Similarly, monthly radio timing ephemerides of the Crab pulsar are available from the Jodrell Bank Centre for Astrophysics (Lyne et al. 1993).¹³

We also investigated the rotation-powered (recycled) millisecond pulsars, PSR J0218+4232 (binary; $P \simeq 2.3$ ms), PSR B1937+21 (isolated; $P \simeq 1.55$ ms) and PSR B1821-24A (isolated; $P \simeq 3.05$ ms). These pulsars have hard non-thermal emission and (very) narrow pulses in their light curves (see e.g., Kuiper & Hermsen 2003; Kuiper et al. 2004; Gotthelf & Bogdanov 2017), providing excellent timing calibration targets for the NuSTAR X-ray bands. However, concurrent NICER and NuSTAR observations only exist for PSR B1821-24A. We used the NICER observations of this pulsar performed during 2017 to generate an accurate timing model and construct a high-statistics NICER pulse profile to compare with the deep

¹³ See <http://www.jb.man.ac.uk/~pulsar/crab.html>.




Figure 7. Comparison of arrival times of few selected pulsars quasi-simultaneously observed by NICER and NuSTAR. NuSTAR data are systematically leading by ~ 5 ms with respect to NICER. Note that the fastest pulsars need more than one full rotation to complete these 5 ms. Also note the slightly higher shift in PSR J2022+3842, that on closer inspection turns out to be an outlier (see Figure 7).

(~ 155 ks) NuSTAR observations performed in April and September 2017.

Finally, we searched for suitable sources among the Accreting millisecond X-ray Pulsars (AMXP) (see e.g., Patruno & Watts 2021, for an overview). Among the 19 AMXPs that went into outburst in the last few years, simultaneous NICER and NuSTAR observations exist for the binary systems IGR J17379-3747 ($P \simeq 2.1$ ms), IGR J17591-2342 ($P \simeq 1.9$ ms), SAX J1808.4-3658 ($P \simeq 2.5$ ms) and Swift J1756.9-2508 ($P \simeq 5.5$ ms; 2 outbursts). Timing models have been constructed for all selected AMXPs from the NICER observations executed during the outburst of each source, which typically lasts a couple of weeks.

4.2.2. Data Reduction

We reduced and analyzed the above pulsar sample according to standard procedures for NuSTAR and NICER data. The photon arrival times at the satellite were corrected to the solar system barycenter (in the TDB reference frame) using the same IDL code in order to avoid any potential differences between software tools. We used the most up-to-date location of the pulsars, the JPL DE200 (Crab) or DE405, and the spacecraft orbital ephemerides. For the selected AMXP sample we corrected the arrival times further for the orbital motion of the ms-pulsar in the binary system.

The results were checked against the standard multi-mission barycenter FTOOL `barycorr` and found to agree at the few microseconds level.

To obtain comparison pulse profiles for each pulsar we selected photons from both instruments in the overlapping 3–10 keV bandpass. Although the different responses of NICER and NuSTAR bias the photons in this interval differently (NICER toward softer photons and NuSTAR toward harder photons), we verified the effect on pulse shapes is negligible for our purposes. We folded the corrected photon arrival times, t , into equal phase bins using same timing model for the both NICER and NuSTAR data, according to




Figure 8. Distribution of time lags between NuSTAR and NICER folded pulse profiles in the 3–10 keV energy band for a sample of eight pulsars measured over 2017 April and 2020 February. The pulse periods for these pulsars range from 1.9 ms to 65.7 ms. Discarding the single outlier at MJD 58133, the weighted average time lag is -4.91 ms, as indicated by a dashed horizontal line. The one-sided 1σ confidence interval is $\sim 68 \mu\text{s}$ (gray area).

the prescription,

$$\phi(t) = \nu \cdot (t - t_0) + \frac{1}{2}\dot{\nu} \cdot (t - t_0)^2 + \frac{1}{6}\ddot{\nu} \cdot (t - t_0)^3. \quad (11)$$

The timing model (ephemeris) in Equation (11) consists of the parameters $(\nu, \dot{\nu}, \ddot{\nu}, t_0)$ representing the spin frequency, first order time derivative and second order time derivative of the frequency at epoch t_0 , respectively.

4.2.3. Results

As shown in Figure 7, NuSTAR data are systematically leading NICER data by ~ 5 ms. To accurately measure this time lag and calibrate the absolute arrival time of the NuSTAR photons we cross-correlate the pulse profile for each pulsar against the reference NICER profile to obtain their relative time lags (Figure 8). The weighted average ΔT is -4.92 ms with a 1σ uncertainty of $124 \mu\text{s}$, or, after removing the PSR J2022+3842 outlier, -4.91 ms with an uncertainty of $68 \mu\text{s}$. The latter is notably compatible with the width of the distribution of the UTC values shown in Figure 5. Apparently, the NuSTAR clock runs about 4.91 ms ahead of the baseline.

5. Clock Correction File Production

The standard NuSTAR clock files released by the SOC via the CALDB prior to version v096 were produced using an IDL interactive script to fit a spline through the raw clock offset measurements from the ground stations. As mentioned previously, due to the rapid temperature-driven drifts of the spacecraft TCXO compared to the relatively sparse ground station passes, the clock correction was only reliable to ~ 2 ms (Madsen et al. 2015).

Since version v096, the distributed clock correction files are produced using the `nustar-clock-utils`,¹⁴ based on the thermal model discussed in this paper. This procedure is automated and a new clock correction file with updated values is generated every ~ 2 weeks. Over the following versions, we

¹⁴ <https://github.com/nustar/nustar-clock-utils>

improved the algorithm and the treatment of missing temperature measurements, bad ground station clock offset measurements, and so on.

The procedure to build the clock file works as follows: (a) we update the latest clock offset measurements from all ground stations, the history of commanded divisor changes, and the temperature measurements; (b) we take note of all time intervals where there are no temperature measurements for more than 10 minutes (e.g., because of reboots of the spacecraft software) (c) for each continuous interval with no missing temperature measurements:

1. The thermal model is applied to the data (as in Figure 2) and the residual clock offsets are calculated with respect to this thermal model (as in Figure 2);
2. An initial model of the residual trends in clock offsets is created by making a robust polynomial interpolation with the outliers removed;
3. The remaining residuals are smoothed using a median filter of 11 clock offset measurements;
4. Finally, the end points of the solution are adjusted to go to zero offset;

(d) for each “bad” interval, having no temperature measurements, we use a straight interpolation of the solution between clock offsets and declare a fixed uncertainty of ~ 1 ms;

The obtained correction has an overall median absolute deviation of $\sim 100 \mu\text{s}$. Finally, (e) we interpolate the correction using the offset and its gradient over a uniform grid of ~ 3 points per NuSTAR orbit; (f) we verify that the solution calculates the interpolation far from grid points with adequate accuracy; (g) we apply a constant offset of 4.91 ms, the best-fit cross-match from Section 4.2; (h) we create a clock correction file that contains four columns: the time at grid points, the clock offsets, the clock offset gradient, and the error bar calculated from the scatter of clock offsets in the days around the grid point.

Each new clock correction file is then tested as follows: we barycenter ~ 45 test data sets obtained by the same number of NuSTAR observations of the Crab, PSR B1821-24A and PSR 1937+21. As periodic calibration observations of the Crab are executed, or any other good pulsar data are available, we add these to our test bench in order to track possible degradations of the clock performance. We calculate TOAs using standard templates aligned with a standard X-ray observation (a previous verified NICER or NuSTAR profile) and verify that they are consistently within $\pm 100 \mu\text{s}$ of the expected arrival time throughout the mission time, with the exception of “problematic” intervals with no Malindi passes, missing temperature information, or other technical problems (see Appendix A).

6. Conclusions

Over the years since the mission launch, we have conducted a deep study of NuSTAR’s timing performance. We found the reference PPS signal produced by the spacecraft’s TCXO oscillator, and used to time tag the events, is temperature-dependent. We characterized this temperature dependence and its change over the course of the mission, probably due to the aging of the TCXO. In this paper, we describe in detail this temperature dependence and how, correcting for it and using pulsar observations for a final alignment to UT time, we can achieve a sub-100 μs timing accuracy. Finally, we describe

how this temperature model is used to produce NuSTAR’s clock correction files distributed with the NuSTAR Calibration database (CALDB). We report on a 4.91 ms offset of the clock offset measured by the ground stations, now accounted for in the clock correction files.

We conclude that, running the barycentering process with the official FTOOL `barycorr` using the clock correction files distributed by the NuSTAR SOC with the CALDB, NuSTAR event times can be trusted to the $\sim 65 \mu\text{s}$ level ($1-\sigma$) throughout the history of the mission, except for a few “bad” intervals whose list is continuously updated in the NuSTAR SOC page. The performance is continuously monitored using an ever-increasing set of test observations of multiple pulsars, in order to single out corner cases where the correction breaks down and/or promptly catch a possible degradation of the timing solution.

M.B. thanks the Fulbright Visiting Scholar Program for supporting a 9 month visit at Caltech. D.J.W. acknowledges financial support from STFC via an Ernest Rutherford fellowship. The authors wish to thank the NuSTAR X-ray Binaries WG, the Magnetars/RPP WG, and the Timing WG, for the hard work and fruitful discussions that led to, and improved, this work; Victoria Kaspi, Jill Burnham, Alessandro Riggio, Andrea Sanna, Paul Ray, Iulia Deneva, Franco Buffa and Alessio Trois for insightful discussions and suggestions. They also wish to thank the referee for their comments, that helped improving the quality of this paper. This research has made use of data and/or software provided by the High Energy Astrophysics Science Archive Research Center (HEASARC), which is a service of the Astrophysics Science Division at NASA/GSFC.

Facilities: NuSTAR, NICER, HEASARC.

Software: Nustar Clock Utils,¹⁵ Stingray (Huppenkothen et al. 2016, 2019), HENDRICS (Bachetti 2018), astropy (Price-Whelan et al. 2018), PINT (Luo et al. 2019), HEASoft/FTOOLS (Blackburn 1995; Blackburn et al. 1999), ATNF pulsar catalog (Manchester et al. 2005), PRESTO (Ransom 2011), corner (Foreman-Mackey 2016), emcee (Foreman-Mackey et al. 2013), IDL,¹⁶ Veusz,¹⁷ Holoviews.¹⁸

Appendix A Example Diagnostic Plots for Clock File Testing

To complement the creation of new clock files, we set up a test bed with an ever-increasing number (currently ~ 45) of sample data sets containing at the moment:

1. Four observations of PSR B1821-24A.
2. One observation each for PSR B1937+21, PSR J2022+3842, IGR J17379-3747.
3. 40 observations of the Crab pulsar.

For the Crab observations, we restricted the number of photons to a manageable level (100,000) for each observation, more than sufficient to precisely measure the alignment with the reference profile. However, dead-time correction would require using all the photons in the observation. Therefore, we created a reference dead-time-affected profile by cross-correlating its dead-time-*corrected* analogous to the NICER

¹⁵ <https://github.com/NuSTAR/nustar-clock-utils>

¹⁶ <https://www.harrisgeospatial.com/Software-Technology/IDL>

¹⁷ <https://veusz.github.io>

¹⁸ <http://holoviews.org>




Figure 9. Application of the absolute timing correction to a sample of four different pulsars with pulse periods 2.1 ms (IGR J17379-3747), 3.05 ms (PSR B1821-24A), 33.5 ms (the Crab), and 48.6 ms (PSR J2022+3842), respectively. NuSTAR data are corrected using clock correction file v105 and v108. All profiles are derived for the 3–10 keV band.




Figure 10. Diagnostic plot from the `nustar-clock-utils` showing the alignment of the Crab pulse over the history of NuSTAR observations. The pulses are not dead-time corrected (hence the distortion of the profile) for performance reasons, but the template pulse was aligned so that its dead-time-corrected version coincided with the NICER profile $\pm 100\mu\text{s}$.

profile once and for all. This allowed us to have a robust dead-time-affected reference profile to be used with fast checks using only 100,000 photons per observation.

Figures 9 and 10 show examples of the diagnostic plots used in the process.

Appendix B Pulsar Searches with NuSTAR: Known Issues

NuSTAR’s detection chain has a high nominal time resolution and, as we demonstrated in this paper, allows an absolute timing precision of better than $\sim 100\ \mu\text{s}$ in most observations. This allows one, in principle, to search for very fast pulsars in NuSTAR data, with frequencies well above the physical limits for NS rotation. However, there are a few known issues that people using NuSTAR for pulsar searches should keep in mind. In this Appendix we will rapidly go through them.















The first, obvious, issue is dead time. For the purpose of this Section, it is sufficient to note that NuSTAR’s 2.5 ms dead time produces a frequency-dependent distortion of the PDF response. This means that pulsar searches will have to use a frequency-dependent detection threshold or try to correct the shape of the PDS. The problem is treated extensively in other papers (e.g., van der Klis 1989; Zhang et al. 1995). The presence of two identical detectors in NuSTAR allows one to partly work around the problem (Bachetti et al. 2015; Bachetti & Huppenkothen 2018).

All NuSTAR observations are executed in *charge pump* mode (CPMODE; Miyasaka et al. 2009). NuSTAR’s detectors accumulate charges continuously due to internal (leakage) currents and other sources of electronic noise as well as the charge deposited by incident X-rays. To prevent the saturation of the readout electronics, a clock-synchronized feedback circuit removes this additional charge every 1.123 ms (spacecraft time) in a “CPMODE reset.” The instrument experiences 30–40 μs of dead time every time this happens.

A high-frequency pulsation search can often “detect” these resets as a strong oscillation close to $\sim 1/1.123\text{s} = 890.5\text{ Hz}$ and/or possible harmonics and aliases of this frequency. The actual frequency can change slightly based on temperature; also, being in spacecraft time, it sometimes goes undetected after barycentering. It is easy to single out the harmonics, as they are exact multiples of the fundamental. The aliases are just slightly more tricky, as they can represent non-obvious reflections of any harmonics about the Nyquist frequency. They can be ruled out by changing the Nyquist frequency itself by using a light curve sampled at a different rate: if the feature does not change frequency, one can safely rule out that the feature is an alias of the 890 Hz frequency.

In very early observations (prior to 2012 August), there was an additional source of periodic dead time from housekeeping operations being run every 1, 4, and 8 s. A PDS of these early observations will typically show strong features at 0.125, 0.25, and all integer frequencies up to $\sim 30\text{ Hz}$.

ORCID iDs

Matteo Bachetti  <https://orcid.org/0000-0002-4576-9337>
 Craig B. Markwardt  <https://orcid.org/0000-0001-9803-3879>
 Brian W. Grefenstette  <https://orcid.org/0000-0002-1984-2932>
 Eric V. Gotthelf  <https://orcid.org/0000-0003-3847-3957>
 Lucien Kuiper  <https://orcid.org/0000-0002-7889-6586>
 Didier Barret  <https://orcid.org/0000-0002-0393-9190>
 Andrew Davis  <https://orcid.org/0000-0002-9922-8915>
 Felix Fürst  <https://orcid.org/0000-0003-0388-0560>
 Karl Forster  <https://orcid.org/0000-0001-5800-5531>
 Fiona A. Harrison  <https://orcid.org/0000-0003-2992-8024>
 Kristin K. Madsen  <https://orcid.org/0000-0003-1252-4891>
 Hiromasa Miyasaka  <https://orcid.org/0000-0002-8074-4186>
 John A. Tomsick  <https://orcid.org/0000-0001-5506-9855>
 Dominic J. Walton  <https://orcid.org/0000-0001-5819-3552>

References

- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2018, *ApJS*, **235**, 37
 Arzoumanian, Z., Gendreau, K. C., Baker, C. L., et al. 2014, *Proc. SPIE*, **9144**, 914420
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Bachetti, M. 2018, HENDRICS: High ENergy Data Reduction Interface from the Command Shell, Astrophysics Source Code Library, ascl:1805.019
 Bachetti, M., Harrison, F. A., Cook, R., et al. 2015, *ApJ*, **800**, 109
 Bachetti, M., & Huppenkothen, D. 2018, *ApJL*, **853**, L21
 Blackburn, J. K. 1995, in ASP Conf. Ser. 77, Astronomical Data Analysis Software and Systems IV (San Francisco, CA: ASP), 367
 Blackburn, J. K., Shaw, R. A., Payne, H. E., & Hayes, J. J. E. 1999, FTOOLS: A general package of software to manipulate FITS files, Astrophysics Source Code Library, ascl:9912.002
 Cusumano, G., La Parola, V., Capalbi, M., et al. 2012, *A&A*, **548**, A28
 Deneva, J. S., Ray, P. S., Lommen, A., et al. 2019, *ApJ*, **874**, 160
 Foreman-Mackey, D. 2016, *JOSS*, **1**, 24
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
 Gotthelf, E. V., & Bogdanov, S. 2017, *ApJ*, **845**, 159
 Harrison, F. A., Craig, W. W., Christensen, F. E., et al. 2013, *ApJ*, **770**, 103
 Huppenkothen, D., Bachetti, M., Stevens, A. L., et al. 2019, *ApJ*, **881**, 39
 Huppenkothen, D., Bachetti, M., Stevens, A. L., Migliari, S., & Balm, P. 2016, Stingray: Spectral-timing software, Astrophysics Source Code Library, ascl:1608.001
 Kuiper, L., & Hermsen, W. 2003, arXiv:astro-ph/0312204
 Kuiper, L., & Hermsen, W. 2009, *A&A*, **501**, 1031
 Kuiper, L., & Hermsen, W. 2015, *MNRAS*, **449**, 3827
 Kuiper, L., Hermsen, W., & Stappers, B. 2004, *AdSpR*, **33**, 507
 Landsberg, P. T. 1955, *JChPh*, **23**, 1079
 Lorimer, D. R. 2008, *LRR*, **11**, 8
 Luo, J., Ransom, S., Demorest, P., et al. 2019, PINT: High-precision pulsar timing analysis package, Astrophysics Source Code Library, ascl:1902.007
 Lyne, A. G., Pritchard, R. S., & Graham Smith, F. 1993, *MNRAS*, **265**, 1003
 Madsen, K. K., Harrison, F. A., Markwardt, C. B., et al. 2015, *ApJ Supp. Ser.*, **220**, 8
 Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *AJ*, **129**, 1993
 Miyasaka, H., Harrison, F. A., Cook, W. R., et al. 2009, *Proc. SPIE*, **7435**, 74350Q
 Newville, M., Stensitzki, T., Allen, D. B., & Ingargiola, A. 2014, LMFIT: Nonlinear Least-Square Minimization and Curve-Fitting for Python, Zenodo, doi:10.5281/zenodo.11813
 Patruno, A., & Watts, A. L. 2021, in Timing Neutron Stars: Pulsations, Oscillations and Explosions, ed. M. Belloni et al. (Berlin: Springer)
 Ransom, S. 2011, PRESTO: PulsAR Exploration and Search Toolkit, Astrophysics Source Code Library, ascl:1107.017
 Saito, Y., Kawai, N., Kamae, T., et al. 1997, *ApJL*, **477**, L37
 Su, W. 1996, in Proc. 1996 IEEE International Frequency Control Symp. (Piscataway, NJ: IEEE), 890
 Takahashi, M., Shibata, S., Torii, K., et al. 2001, *ApJ*, **554**, 316
 Taylor, J. H. 1992, *RSPTA*, **341**, 117
 van der Klis, M. 1989, in Timing Neutron Stars: Proceedings of the NATO Advanced Study Institute on Timing Neutron Stars, ed. H. Ögelman & E. P. J. van den Heuvel, Vol. 4 (New York: Kluwer Academic), 27
 van der Klis, M. 2006, in Compact Stellar X-Ray Sources, ed. W. Lewin & M. van der Klis (Cambridge: Cambridge Univ. Press), 39
 Vig, J., & Meeker, T. 1991, in Proc. 45th Annual Symp. on Frequency Control (Piscataway, NJ: IEEE), 77
 Zhang, W., Jahoda, K., Swank, J. H., Morgan, E. H., & Giles, A. B. 1995, *ApJ*, **449**, 930