

Pulse profile modelling of the 2024 outburst of the accreting millisecond pulsar SRGA J144459.2–604207

Bas Dorsman¹,¹★ Tuomo Salmi²,²★ Anna L. Watts³,^{1,3}★ Mason Ng^{4,5}, Anna Bobrikova⁶,
Alessandro Di Marco⁷, Duncan K. Galloway^{8,9}, Sebastien Guillot¹⁰, Mariska Hoogkamer¹,
Yves Kini¹¹, Fabio La Monaca¹², Vladislav Loktev¹², Matteo Lucchini¹, Christian Malacaria¹³,
Ying-Han Mao^{14,15}, Alessandro Papitto¹³ and Juri Poutanen⁶

¹Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, NL-1098XH Amsterdam, the Netherlands

²Department of Physics, University of Helsinki, PO Box 64, FI-00014 Helsinki, Finland

³Gravitation and Astroparticle Physics Amsterdam (GRAPPA), University of Amsterdam, Science Park 904, NL-1098XH Amsterdam, the Netherlands

⁴Department of Physics, McGill University, 3600 rue University, Montréal, QC H3A 2T8, Canada

⁵Trottier Space Institute, McGill University, 3550 rue University, Montréal, QC H3A 2A7, Canada

⁶Department of Physics and Astronomy, University of Turku, Turku, FI-20014, Finland

⁷INAF Istituto di Astrofisica e Planetologia Spaziali, Via del Fosso del Cavaliere 100, I-00133 Roma, Italy

⁸School of Physics and Astronomy, Monash University, VIC 3800, Clayton, Australia

⁹Institute for Globally Distributed Open Research and Education (IGDORE)

¹⁰IRAP, CNRS, 9 avenue du Colonel Roche, BP 44346, F-31028; Université de Toulouse, CNES, UPS-OMP, F-31028, Toulouse, France

¹¹Dipartimento di Fisica, Università degli Studi di Roma ‘Tor Vergata’, Via della Ricerca Scientifica 1, I-00133 Roma, Italy

¹²School of Mathematics, Statistics, and Physics, Newcastle University, Newcastle upon Tyne NE1 7RU, UK

¹³INAF Osservatorio Astronomico di Roma, Via Frascati 33, I-00078 Monte Porzio Catone (RM), Italy

¹⁴School of Astronomy and Space Science, Nanjing University, Nanjing 210023, P R China

¹⁵Key Laboratory of Modern Astronomy and Astrophysics (Nanjing University), Ministry of Education, Nanjing 210023, P R China

Accepted 2026 July 10. Received 2026 July 9; in original form 2026 May 18

ABSTRACT

Pulse profile modelling via relativistic ray-tracing can constrain the system parameters of neutron stars, notably their mass and radius. Among these objects, accreting millisecond pulsars (AMPs) are promising targets, because they are bright in X-rays and their potentially polarized radiation can lead to complementary constraints on the emission geometry. We perform combined analysis of NICER (Neutron star Interior Composition Explorer) and *IXPE* (*Imaging X-ray Polarimetry Explorer*) observations of the recently discovered the 448-Hz pulsar SRGA J144459.2–604207, with *IXPE* providing X-ray polarization information. NICER and *IXPE* jointly favour a large mass and radius for our best-fitting model, for which the neutron star has two independent hotspots. The primary hotspot is centred near the Northern rotational pole, the secondary in the Southern hemisphere, and the observer inclination is in the range 50°–75°. The primary hotspot is large (up to half the surface area) and contributes the majority of the non-pulsed X-rays, while the secondary is hotter and the major contributor to the overall pulse profile shape. However, many parameters are inferred to be near the prior bounds, which could indicate that the model does not adequately account for important physics. Furthermore, we tested several different methodologies for joint analysis of the two data sets: the results are sensitive to the method used, something that merits further study with synthetic data. In the future, we expect simultaneously recorded data will lead to improved parameter constraints, especially when multiband and polarized data are combined.

Key words: accretion, accretion discs – equation of state – stars: individual: SRGA J144459.2–604207 – stars: neutron – X-rays: binaries.

1 INTRODUCTION

Neutron stars (NSs) are rich astrophysical laboratories for the study of the equation of state (EoS) of ultradense nuclear matter,

under conditions that cannot be replicated on Earth. Both the particle types (e.g. neutrons, hyperons, and deconfined quarks) present in NS cores, as well as their interactions, remain uncertain despite intense theoretical, experimental, and astrophysical study (see e.g. K. Hebeler et al. 2015; L. Tolos & L. Fabbietti 2020; J. M. Lattimer 2021; K. Chatziioannou et al. 2025; H. Koehn et al. 2025).

One way to constrain the EoS is to measure the NS mass (M) and equatorial radius (R_{eq}), since the M – R_{eq} relation depends

* E-mail: b.dorsman@uva.nl (BD); tuomo.salmi@helsinki.fi (TS); A.L.Watts@uva.nl (ALW)

on the EoS via the stellar structure equations. One can do this measurement using pulse profile modelling (PPM) of millisecond X-ray pulsars, which are fast-rotating NSs with X-ray emitting hotspots. PPM is a relativistic ray-tracing technique used to estimate stellar parameters, such as M and R_{eq} , by modelling the observed energy-phase dependent pulse profiles (see e.g. K. R. Pechenick, C. Ftacals & J. M. Cohen 1983; M. C. Miller & F. K. Lamb 1998; J. Poutanen & M. Gierliński 2003; S. M. Morsink et al. 2007; S. Bogdanov et al. 2019, 2021, and references therein). The technique also constrains the properties of the hotspots, such as their size, location, and temperature.

So far, the technique has been most successfully applied to rotation powered millisecond pulsars (RMPs), for which the hotspots are thought to be generated at the magnetic poles due to heating by magnetospheric return currents (see e.g. A. K. Harding & A. G. Muslimov 2001). PPM has been applied to data from the Neutron star Interior Composition Explorer (NICER; K. C. Gendreau et al. 2016), often supplemented by data from the *X-ray Multi-Mirror Newton (XMM-Newton)* mission to help constrain the non-source background. Current constraints on M and R_{eq} are at the $\sim \pm 10$ per cent level (68 per cent credible interval, CI) and hotspot maps suggest that magnetic fields are not centred dipoles (A. V. Bilous et al. 2019; M. C. Miller et al. 2019, 2021, 2026; T. E. Riley et al. 2019, 2021; T. Salmi et al. 2022, 2023, 2024a, b; D. Choudhury et al. 2024; A. J. Dittmann et al. 2024; S. Vinciguerra et al. 2024; M. Hoogkamer et al. 2025; L. Mauviard et al. 2025; L. Qi et al. 2025; Y. Kini et al. 2026). RMPs lend themselves well to PPM, despite low X-ray fluxes, because their X-ray pulse profiles are stable over a long period of time, such that pulse profiles can be constructed from years of observational data. For several of these RMPs, radio timing provides informative priors on M , improving the overall constraints (see e.g. H. T. Cromartie et al. 2020; E. Fonseca et al. 2021; D. J. Reardon et al. 2024).

PPM can also be applied to pulsations from rapidly rotating accreting NSs in low-mass X-ray binaries (LMXBs). In such systems, when the low-mass companion fills its Roche lobe, matter is transferred through the inner Lagrange point, and owing to its specific angular momentum, the transferred material forms an accretion disc around the NS. During outburst episodes (J.-P. Lasota 2001; C. O. Heinke et al. 2025), material from the inner edge of the disc is accreted onto the surface of the NS, giving rise to X-ray emission. PPM can be applied to two types of pulsating NSs in LMXBs. The first type are thermonuclear burst oscillation (TBO) sources. In that case, during Type I (thermonuclear) X-ray bursts, surface anisotropies on the rotating NS appear, causing oscillations to be observed (A. L. Watts 2012; S. Bhattacharyya 2022). Performing PPM for this class of pulsations has so far proved challenging due to pulse variability and computational cost (Y. Kini et al. 2023, 2024a, b, 2025). The second type are the persistent pulsations of accreting millisecond pulsars (AMPs), where the magnetic field of the NS channels the accreted matter towards the magnetic poles of the NS, leading to localized surface heating around the magnetic poles (A. Patruno & A. L. Watts 2021; T. Di Salvo & A. Sanna 2022). For observers at favourable viewing angles, this leads to the detection of X-ray pulsations: 27 AMPs are now known. Due to Comptonization of seed photons by hot electrons inside the in-falling material, the radiation should become linearly polarized (K. Viironen & J. Poutanen 2004).

AMPs are interesting sources for PPM study. First, because they are brighter (hundreds of counts per second, compared to less than 1 count per second for RMPs, see e.g. P. Bult et al. 2019; T.

Salmi et al. 2022), they have the potential of providing tight constraints. Furthermore, due to their diverse phenomenology, other independent and complementary methods can be applied, permitting cross-verification of potential systematic errors in methods (A. L. Watts et al. 2016; A. Li et al. 2025). Comparison to the results obtained for RMPs should also help us better understand the NS population and formation history, because AMPs are believed to be the progenitors of RMPs (although there are some caveats to this, F. D’Antona & M. Tailo 2022). PPM of AMPs can also place constraints on the accretion flow phenomena, such as the accretion disc and accretion column (J. Poutanen, A. Ibragimov & M. Annala 2009; V. Ahlberg, J. Poutanen & T. Salmi 2024; B. Dorsman et al. 2025).

However, PPM of AMPs also introduces new challenges. AMPs only sometimes go into outburst and their variability forces the selection of shorter data sets compared to RMPs (enough photons can typically still be gathered because AMPs are more luminous). There are also other sources of X-ray emission beyond just the surface hotspots. These include the inner disc region, as well as the accretion column. These components need to be taken into account in the PPM, and may lead to degeneracies in the modelling (B. Dorsman et al. 2025, 2026). However, X-ray polarization data – as recently provided by the *Imaging X-ray Polarimetry Explorer* observatory (*IXPE*; M. C. Weisskopf et al. 2022; A. Papitto et al. 2025) – have been predicted to provide additional constraints that can alleviate this issue, because polarization provides an independent measure of the geometry (K. Viironen & J. Poutanen 2004; T. Salmi et al. 2021; A. Bobrikova et al. 2023; T. Salmi et al. 2025).

The first PPM studies for AMPs, which laid much of the theoretical groundwork, made use of data from the *Rossi X-ray Timing Explorer* (J. Poutanen & M. Gierliński 2003; M. Gierliński & J. Poutanen 2005; D. A. Leahy, S. M. Morsink & C. Cadeau 2008; D. A. Leahy et al. 2009; S. M. Morsink & D. A. Leahy 2011; T. Salmi, J. Nättälä & J. Poutanen 2018). More recently, B. Dorsman et al. (2026) carried out PPM using NICER data for the AMP SAX J1808.4–3658 from its 2019 and 2022 outbursts, finding a wide range of allowed parameters depending on how the accretion disc is modelled. They concluded that, in the soft X-ray band of NICER (at least for this source), there exists a degeneracy between potential sources of the non-pulsed radiation component: radiation can come either from the disc or from the NS surface, if part of the hotspot is continuously visible throughout a full rotation. Unfortunately, *IXPE* was unable to observe SAX J1808.4–3658 during the 2022 outburst – so there were no polarization data to break the degeneracy.¹

In 2024 February, however, S. V. Molkov et al. (2024) reported the discovery of a new 447.9-Hz AMP, SRGA J144459.2–604207. The AMP nature was confirmed using NICER (M. Ng et al. 2024), and shortly afterwards, A. Papitto et al. (2025) reported the discovery of polarized emission using data from *IXPE*. This makes SRGA J144459.2–604207 the first AMP for which polarized pulse profiles are available. This opens up the possibility of doing joint PPM for both data sets, in the hope of obtaining tighter constraints on the system parameters. This is the main goal of this paper.

While previous PPM studies have combined pulsed and unpulsed (phase-averaged) data sets, this is the first PPM study to

¹There are *IXPE* observations of SAX J1808.4–3658 from 2025 that are currently being analysed (C. Ballocco et al. 2025).

jointly analyse two pulsed data sets. Our secondary aim is thus to learn lessons for the future joint analysis of multiple data sets in the context of PPM of AMPs. This topic is of major importance for the upcoming *enhanced X-ray Timing and Polarimetry (eXTP)* mission (A. Li et al. 2025; S.-N. Zhang et al. 2025), which has spectral and polarimetry instruments on board that can observe simultaneously. We note that the effect on parameter constraints of combining non-polarized X-ray data with polarization data were previously investigated with simulated *eXTP* data by A. Li et al. (2025).

This paper is structured as follows. Section 2 describes the PPM methodology applied here. Section 3 describes the preparation of the NICER and *IXPE* data. Section 4 describes the results of the separate as well as joint PPM analyses. Section 5 discusses the results, and provides recommendations for future work. Finally, we conclude in Section 6.

2 METHODOLOGY

To perform PPM, we make use of the open source X-RAY PULSE SIMULATION AND INFERENCE software (X-PSI, v3.2.0, T. E. Riley et al. 2023), used extensively for the analysis of RMPs and adapted for AMP studies including polarimetric data (B. Dorsman et al. 2025, 2026; T. Salmi et al. 2025). We consider here two data sets, observed by NICER and *IXPE*, with the aim of improving the accuracy and precision of parameter inferences. For context, the two observations took place separated by a gap of 3.6 d (see Section 3 for more detail). The outburst flux and pulse profiles evolved between (and during) the observations. This variation implies that the parameters that govern the hotspots and accretion disc, for example, may change significantly between the two observations and we thus consider these ‘variable’. On the other hand, a small subset of the parameters are considered ‘constant’: M , R_{eq} , distance (D), and cosine of the observer inclination ($\cos i$), because they must be unchanging on the timescale of one outburst. We also initially assume hydrogen column density (N_{H}) to be constant, but locally the effective X-ray absorption could evolve as the outburst proceeds, so we also consider analyses of the data where we assume this parameter to be variable. There are multiple Bayesian ways to analyse two sets of data, each with drawbacks and benefits, and here we will attempt three approaches. We will briefly introduce them now, but describe these in more detail in Section 2.2.

First, we analyse the two data sets independently. This approach is computationally feasible and a useful baseline, giving two independent sets of results that can be cross-verified. We also combine the independent posterior probability distributions after the fact, leading to joint posteriors. Second, we carry out a two-step Bayesian analysis where the posteriors from analysis of one data set are used as ‘updated’ priors for the other (see Section 2.2 for a more detailed explanation). The prior update only takes place for any constant parameters; variable parameters are free to change and are thus not informed by a previous data set. This second strategy is also computationally affordable. In this paper, we only update the prior density functions separately for each parameter using 1D cumulative distribution functions, noting that this method may result in a reduction of accuracy due to the loss of higher dimensional shape. Third, we construct a joint model to analyse the two data sets together, carrying out one nested sampling analysis with the joint model to infer parameters of both data sets simultaneously. Variable parameters will have two instances, one for each data set, meaning that the number of

parameters governing this model is significantly increased. This method improves over the previous analysis in that it overcomes the loss of higher dimensional shape, but the increased number of parameters represents a considerable computational challenge. The remainder of this section describes the physical model deployed here, the parameter inference process, and finally, the prior probability distributions and the reasoning behind each.

2.1 Physical model

This section describes the model that produces the pulse profiles. This model is mostly the same as in B. Dorsman et al. (2025, 2026) and consists of two parts: the NS and the accretion disc. Let us first describe the NS: the model that describes the X-ray emission of the surface hotspots. The underlying physics that governs the NS surface emission has been studied over decades (K. R. Pechenick et al. 1983; H. Riffert & P. Meszaros 1988; M. C. Miller & F. K. Lamb 1998; J. Poutanen & M. Gierliński 2003). We give a short overview below, but for a comprehensive description we recommend S. Bogdanov et al. (2019). The paths of light rays that radiate off the NS surface are computed in the oblate Schwarzschild + Doppler approximation, which assumes an oblate star embedded in a Schwarzschild space-time (S. M. Morsink et al. 2007; M. AlGendy & S. M. Morsink 2014). Corrections are made to account for general relativistic light bending, light path time delay, emission angle aberration, and relativistic Doppler boosting. This approximation is computationally efficient and sufficiently accurate for millisecond pulsars that rotate at a frequency below ~ 600 Hz (see e.g. G. A. Oliva & F. Frutos-Alfaro 2021; H. O. Silva et al. 2021; Z. Jakab & S. M. Morsink 2025). The radiating hotspots on the NS surface are parametrized by spherical caps with circular perimeters. The colatitudinal position is defined by $\theta_{\text{p/s}}$ and phase $\phi_{\text{p/s}}$ defines the azimuthal position of the centres of the primary/secondary hotspots. The angular radius of the hotspots is defined by $\zeta_{\text{p/s}}$. Their atmospheric radiation pattern is given by the Comptonization model of A. Bobrikova et al. (2023) which includes polarization, and is governed by three parameters (omitting here the p/s subscript to simplify notation): seed photon temperature (T_{seed}), electron temperature (T_{e}), and Thomson optical depth (τ). The atmospheric parameters are constant across the surface of the hotspots. We call the model of an NS with a single hotspot of uniform distribution ST (for single temperature), ST-S in the case of two hotspots that are antipodal twins with shared atmospheric parameters, and ST-U in the case of two independent (unshared) but non-overlapping hotspots. The hotspot emission bound for the observer is calculated at each rotational step of the NS and collected in the appropriate phase and energy bin of the simulated pulse profile depending on energy shift and light path.

The second part of the model is the accretion disc, which was first implemented into X-PSI by B. Dorsman et al. (2025). B. Dorsman et al. (2025) only implemented phase-independent disc emission (i.e. `diskbb`, see also K. Mitsuda et al. 1984; K. Makishima et al. 1986), but the presence of a disc naturally implies the geometric eclipse of NS radiation by the disc. For this reason, we extend the model used in B. Dorsman et al. (2025, 2026) to include disc occultation. This model is used in the analysis of the NICER data, but not in the final analysis of *IXPE* data, as explained in Section 4.2. This addition is also motivated by the findings of S. V. Molkov et al. (2024), who required partial hotspot occultation by the disc to reproduce the pulse profiles of SRGA J144459.2–604207 observed by the Astronomical Roent-

gen Telescope X-ray Concentrator (ART-XC) instrument on board the *Spectrum–Roentgen–Gamma* observatory (SRG; R. Sunyaev et al. 2021). To assess whether a photon trajectory is intercepted by the disc, we follow the approximate light-bending treatment in appendix C of A. Ibragimov & J. Poutanen (2009). For each ray, we compute the radius at which it crosses the disc plane: if this radius is smaller than the inner disc radius (R_{in}), the ray remains unobscured; otherwise, it is blocked. This typically affects mostly the secondary hotspot and does not yet include disc blocking of higher order images, although that contribution is only relevant for very compact stars. This disc occultation module is implemented in X-PSI from v3.2.0 and discussed in Y.-H. Mao et al. (2026). Afterwards, the non-eclipsed emission from the hotspots is added to the emission of the disc. The final pulse profile is obtained after attenuating the combined NS and disc flux by interstellar absorption and convolving through the response of the observing telescope.

2.2 Parameter inference

The main goal of this work is to estimate the posteriors of model parameters. This section presents a short review of the Bayesian inference methodology relevant for this work, followed by an description of the strategies explored to analyse the two data sets. The posterior is expressed with Bayes’ theorem as follows:

$$P(\theta|D) = \frac{\mathcal{L}(D|\theta)\pi(\theta)}{P(D)}. \quad (1)$$

Here, $P(\theta|D)$, the posterior, is the probability of the set of parameters θ in model $M(\theta)$ given data set D . The likelihood $\mathcal{L}(D|\theta)$ is the probability of obtaining data set D given θ . $\pi(\theta)$ is the prior, the probability of θ before considering D . The model $M(\theta)$ and parameters θ were described in Section 2.1 and the priors are described in Section 2.3. An overview of the parameters and their priors is given in Table 1. A common way to think about equation (1) is that it prescribes how prior knowledge of model parameters is ‘updated’ upon the consideration of D , thus defining the posterior.

$P(D)$, often denoted with the symbol $\mathcal{Z}$, is the probability of obtaining D , and is known as the marginal likelihood or evidence. Notice that $\mathcal{Z}$ does not depend on θ , and therefore only serves as a normalization factor for the posterior. When comparing two or more models however, $\mathcal{Z}$ is crucial because the ratio of two evidences defines the probability ratio of either model explaining the data, also known as the Bayes factor.

To estimate the posterior, we use nested sampling (J. Skilling 2004). Nested sampling in Bayesian parameter inference is an algorithm that explores the parameter space more thoroughly and efficiently than ordinary random sampling. The top level idea is that the parameter space is mapped step by step using nested hypervolumes of prior space. The algorithm starts out by drawing samples from the full prior space but moves inward towards higher likelihood while recording an estimated prior probability ‘mass’ that was removed at each step. Upon reaching some user-defined stopping criterion, one has obtained an estimate of the likelihood surface across parameter space, which is then used to estimate the posteriors and the overall $\mathcal{Z}$. We use the MULTINEST software (F. Feroz, M. P. Hobson & M. Bridges 2009; J. Buchner et al. 2014) to perform this task. In our case, the stopping criterion is reached when the estimated possible remaining contribution from the current set of live points to the evidence $\Delta \ln(\mathcal{Z})$ is less than 0.5 of the currently estimated evidence $\ln(\mathcal{Z})$

Table 1. Overview of model parameters and their prior distributions. Some prior distributions are made up of two parts: the density distribution and bounds. In those cases, the densities describe the shape of the prior, such as a normal distribution, but the cut-offs define hard lower and upper limits, denoted as (lower, upper). $U(\text{lower, upper})$ is a uniform distribution with lower and upper bounds. $N(\mu, \sigma)$ is a normal distribution with mean μ and a standard deviation σ . The subscripts p/s indicate the primary and secondary hotspot, because these parameters exist for both hotspots. The bottom part of the table lists the α -parameters, for which the priors change between the analyses of *IXPE* data or the joint analysis (combination method 3).

Parameter (unit)	Description	Prior density
Pulsar		
M ($M_{\odot}$)	Mass	$U(1, 2.2)^{a,b}$
R_{eq} (km)	Equatorial radius	$U(8, 14)^{a,b}$
D (kpc)	Distance	$N(8.5, 2)^c$
$\cos i$ (–)	Cosine inclination	$U(0, 0.64)$
f (Hz)	Pulsar frequency	447.871561
N_{H} (10^{22} cm^{-2})	Column density	$U(1.9, 2.9)$
χ_0 (deg)	Spin axis position angle	$U(0, 360)$
Hotspots		
$\phi_{\text{p/s}}$ (cycles)	Phase	$U(-0.25, 0.75)$
$\cos \theta_{\text{p/s}}$ (–)	Cosine colatitude	$U(-1, 1)$
$\zeta_{\text{p/s}}$ (deg)	Angular radius	$U(0, 90)$
$T_{\text{seed,p/s}}$ (keV)	Seed photon temperature	$U(0.5, 1.5)$
$T_{\text{e,p/s}}$ (keV)	Electron slab temperature	$U(20, 100)$
$\tau_{\text{p/s}}$ (–)	Thomson optical depth	$U(0.5, 3.5)$
Disc		
T_{in} (keV)	Inner disc temperature	$U(0.01, 0.6)$
R_{in} (km)	Inner disc radius	$U(R_{\text{eq}}, R_{\text{co}})^b$
α -parameters		
α_{N}	Scaling of NICER response	Fixed to 1 ^d $U(0.8, 1.2)^e$
α_{DU1}	Scaling of <i>IXPE</i> DU1 response	$U(0.8, 1.2)^d$ Fixed to 1 ^e
$\alpha_{\text{DU2, DU3}}$	Scaling of <i>IXPE</i> DU2/3 response	$U(0.8, 1.2)^d$ $U(0.95, 1.05)^e$

Notes. ^a These priors are also cut-off by a compactness limit. See the text in Section 2.3 for more detail.

^b The priors of M , R_{eq} , and R_{in} depend on each other because samples are rejected if $R_{\text{in}} > R_{\text{co}} \propto M^{1/3}$ and $R_{\text{in}} < R_{\text{eq}}$. See the text in Section 2.3 for more detail.

^c This prior is truncated below 2.5 kpc and above 10.6 kpc.

^d Individual analyses (when applicable), plus Combination methods 1 and 2.

^e Combination method 3.

(explained in more detail in F. Feroz et al. 2009, 2019). This condition is no guarantee for convergence, and here we also test for convergence by increasing the number of live points until we find that posterior distributions no longer shift – at least within the computational budget.

Now, we describe the strategies employed to separately and jointly analyse the two data sets. Analysis with a single instrument, as we will do with the individual NICER and *IXPE* data sets, is the most straightforward parameter inference scenario. The key ingredients are a function to sample parameters according to the prior, and a function that returns the $\ln$ -likelihood given some parameters and a data set. The case of NICER is simpler, for which the data set consists of a singular pulse profile consisting of X-ray counts. In case of *IXPE*, we have nine different data sets: Stokes I , Q , and U parameters for three different detector units

(DU_{1,2,3}). The total ln-likelihood is the sum over the likelihoods of the individual data sets. For the NICER and *IXPE* Stokes *I* data, we use the Poisson likelihood because the X-ray counting is subject to Poisson noise. For *IXPE* Stokes *Q* and *U*, we use a Gaussian likelihood, which we discuss further in Section 3.2. With these ingredients in hand, $P(\theta|D) \propto \mathcal{L}(D|\theta)\pi(\theta)$ can be estimated through the nested sampling process.

In this work, we are also interested in doing a joint analysis of these two data sets. We describe here three different methods. Combination method 1, combining independent posteriors, is to estimate the joint posterior by combining the results of the separate analyses of the *IXPE* data (D_I) and NICER data (D_N). To do so, we use a similar approach as the one presented in Y. Kini et al. (2024b, see their sect. 2.2.4). To take D_N as an example, the ln-likelihood surface is given by $\ln \mathcal{L}(D_N|\theta_c, \theta_N) = \ln P(\theta_c, \theta_N|D_N) - \ln \pi(\theta_c, \theta_N)$, where we have separated θ_c , the subset of parameters that are constant between observations, and θ_N , the subset of variable parameters (defined in the beginning of Section 2), as observed by NICER. In the case of *IXPE*, these are θ_I .

We then approximate the joint ln-likelihood surface by adding both ln-likelihood surfaces. We do so by using kernel density estimations (KDEs) of the posterior weighted samples and samples drawn from the priors. The joint posteriors are obtained by nested sampling again, now utilizing the KDE-based joint ln-likelihood and the combined prior $\pi(\theta_c, \theta_I, \theta_N)$. The dimensionality of the final nested sampling step can be reduced by marginalizing out the subset of θ_I and θ_N that are only nuisance parameters [not affecting the constant parameters (θ_c)]. It turns out that only R_{in} is not a nuisance parameter. That is because the joint prior, $\pi(\theta_c, R_{in}) \neq \pi(\theta_c)\pi(R_{in})$, because $\pi(R_{in}) = f(M, R_{eq})$, thus causing prior correlation (see Section 2.3 for more detail on each prior). Still, the dimensionality of the joint sampling problem is greatly reduced to at most only six parameters: θ_c (five parameters) and $R_{in,N}$. Because of the small dimensionality and the inexpensive ln-likelihood evaluation (from a KDE) the final sampling step is computationally cheap and fast in real time.

Combination method 2, two-stage Bayesian analysis, is to do parameter inference in two steps. A first round of nested sampling estimates, for the sake of example, D_N : $P(\theta_c, \theta_N|D_N) \propto \mathcal{L}(D_N|\theta_c, \theta_N)\pi(\theta_c, \theta_N)$. We then define $\pi_{\text{updated}}(\theta_c) = P(\theta_c|D_N)$, and $\pi_{\text{updated}}(\theta_c, \theta_I) = \pi_{\text{updated}}(\theta_c)\pi(\theta_I)$. In the case where D_N was analysed first, R_{in} is not a parameter in the subsequent *IXPE* analysis, because no disc is used there. In the reverse case, if D_I was analysed first R_{in} exists and will be drawn using the updated prior distribution of M and R_{eq} . In the NICER-first case $P(\theta_c, \theta_I|D_I, D_N) \propto \mathcal{L}(D_I|\theta_c, \theta_I)\pi_{\text{updated}}(\theta_c, \theta_I)$ is estimated during a second round of nested sampling that considers D_I . The priors of any parameters outside of θ_c are not updated because they are variable and thus not informed by the other data set. In this example, D_N was analysed first, but the order of data set operations could be reversed and (in an idealized world) the analysis leads to equivalent results. However, in practice, a difference could occur due to insufficient exploration of the parameter space during sampling. On top of that, we are only able to update $\pi(\theta_c)$ with separate 1D priors, leading to the loss of any higher dimensional correlation information.

Combination method 3, simultaneous fit, is to use a joint model that simultaneously fits two data sets D_I and D_N , and fit parameters in one round of nested sampling. To do so, a joint

posterior is required, and it is given by

$$P(\theta_c, \theta_I, \theta_N|D_I, D_N) \propto \mathcal{L}(D_I, D_N|\theta_c, \theta_I, \theta_N)\pi(\theta_c, \theta_I, \theta_N). \quad (2)$$

The observational data are independent of each other, so the joint likelihood factorizes as

$$\mathcal{L}(D_I, D_N|\theta_c, \theta_I, \theta_N) = \mathcal{L}_I(D_I|\theta_c, \theta_I)\mathcal{L}_N(D_N|\theta_c, \theta_N), \quad (3)$$

so in ln-space, the joint ln-likelihood is the sum of two ln-likelihoods: $\ln(\mathcal{L}) = \ln(\mathcal{L}_I) + \ln(\mathcal{L}_N)$, one for each data set.

In this method, there exist two instances of those parameters that can vary between the disjoint data sets, one for each data set. This enlarges the parameter space and makes the sampling problem harder. Because of this, and because the emission now needs to be calculated twice (one pulse profile for each data set), this run is more costly than methods 1 and 2. When we attempted to use this method, because of the computational expense, we could only use the limited (antipodal) ST-S model, which led to worse fits. For completeness, we detailed the usage and results of this combination method in Appendix A. The main body of this work uses combination methods 1 and 2, which allowed the usage of the ST-U model, leading to better evidence and better fits.

2.3 Prior distributions

Following the discovery of SRGA J144459.2–604207, a number of studies have used the data from the 2024 outburst to infer parameters, providing helpful prior information for this study. This section describes the prior distributions we adopt and Table 1 provides an overview.

Here, we use a joint flat M - R_{eq} prior, which has been common in preceding PPM analyses (T. E. Riley, G. Raaijmakers & A. L. Watts 2018). However, the prior was wide, typically $M \sim U(1, 3)M_\odot$ and $R_{eq} \sim U(3R_g(1), 16)$, where U is a flat distribution between a lower and upper limit, and $R_g(1)$ is the gravitational radius (GM/c^2) of $1 M_\odot$. While this choice of prior is conservative, it produces many samples that are disfavoured by contemporary EoS theory, experiments and astrophysical results (see e.g. N. Rutherford et al. 2024). Instead, we use a joint flat prior with refined upper and lower limits derived from the EoS-informed prior based on the piecewise-polytropic (PP) model described in N. Rutherford et al. (2024). The prior for M is $U(1, 2.2)M_\odot$ and R_{eq} is $U(8, 14)$ km. These bounds cover almost the whole prior distribution but exclude 1.4 per cent of the prior that has $M > 2.2M_\odot$ and 2.3 per cent of the prior overall. We do not yet adopt the M - R density distribution associated with N. Rutherford et al. (2024), because at the time of the analysis it was not yet possible in X-PSI to use multidimensional priors of arbitrary shape. However, we do note that this would be the preferred method.

Furthermore, compactness limits were applied as usual. Initially, we reject samples beyond the causality limit given by $R_{pole}/R_g(M) > 2.9$ (see e.g. S. Gandolfi, J. Carlson & S. Reddy 2012), where R_{pole} is the polar radius. However, for combination method 2, we increased the lower limit on radius of $R_{pole}/R_g(M) > 3.01$.²

²We encountered numerical issues in X-PSI when the stellar radius lies within its photon sphere (see <https://github.com/xpsi-group/xpsi/issues/633>). We therefore adopted the limit commonly used in NICER RMP analyses. This choice appears justified, as the final posteriors do not approach this regime.

Constraints on the distance to the source come from the measurement of peak flux from the X-ray bursts observed with different instruments, coupled with assessment of the presence or absence of photospheric radius expansion (PRE) in the time-resolved spectroscopic measurements. M. Ng et al. (2024) and C. Malacaria et al. (2025) derived upper limits of 10.6 and 11.5 kpc, respectively, based on the lack of PRE inferred from the bursts observed with NICER and *XMM-Newton*/*NuSTAR*. We adopt the strictest upper limit of the two at 10.6 kpc. We combine this the upper limit with the distance range estimated by S. V. Molkov et al. (2024) of 8–9 kpc, based on more detailed spectral fits during bursts observed with *SRG/ART-XC*. Being conservative, we set up a normal distribution that is centred at the middle of this range, $\mu = 8.5$, but is broader with $\sigma = 2$ kpc, and a lower limit cut-off at 3σ , corresponding to a lower limit of 2.5 kpc. We also note that, T. Fu et al. (2025) estimate a distance of 10.03 ± 0.71 kpc, concluding that roughly 25 per cent of the bursts observed with *Insight-Hard X-ray Modulation Telescope (HMXT)* exhibited PRE. However, we do not find the case for PRE in those bursts convincing, as the time-resolved spectroscopic histories lack a significant radius maximum and contemporaneous temperature minimum (see e.g. D. K. Galloway et al. 2020).

For the observer inclination of the system, a breadth of values have so far been inferred. A. Papitto et al. (2025) fitted the observed pulsations by *IXPE* to a simple NS model. This was based on the analytical relations by J. Poutanen & A. M. Beloborodov (2006) and J. Poutanen (2020) with simplifying assumptions including a spherical NS (no oblateness), constant anisotropy of the emitted radiation and an energy-independent polarization. This model used fixed $M = 1.4M_{\odot}$ and $R = 10$ km, and point-like antiphased hotspots. Using this model they fit Stokes I , Q , and U and infer $i = 74.1^{+5.8}_{-6.3}$. C. Malacaria et al. (2025) perform spectral modelling on *XMM-Newton* and *NuSTAR* data and find $i = 85^{+50}_{-14}$ or 53° (depending on the approach taken to fit disc reflection). Finally, M. Mandal, S. Naik & B. Chhotaray (2025) fit the persistent broadband spectra using NICER and *NuSTAR* data and find $i = 50.3^{+2.0}_{-1.3}$. We use a broad prior from 50° – 90° , which is flat in $\cos i$, covering the breadth of the values that were inferred so far.

We fix the pulsar spin frequency f at 447.871561 Hz, as obtained by M. Ng et al. (2024). For N_H , a range of values has been inferred, using different models. M. Ng et al. (2024) found $(2.90 \pm 0.03) \times 10^{22} \text{ cm}^{-2}$, based on a joint spectral fit of three observation IDs (OBSIDs) of NICER data. C. Malacaria et al. (2025) found $(1.92^{+0.03}_{-0.01}) \times 10^{22} \text{ cm}^{-2}$ in their spectral fitting, and M. Mandal et al. (2025) found $(2.5 \pm 0.1) \times 10^{22} \text{ cm}^{-2}$. Z. Li et al. (2025) analysed the X-ray broad-band spectrum using NICER, *Einstein Probe*, *IXPE*, *Insight-HXMT*, and *INTEGRAL* data. They found $(2.63 \pm 0.05) \times 10^{22} \text{ cm}^{-2}$. We use a broad flat prior of $U(1.9, 2.9) \times 10^{22} \text{ cm}^{-2}$, which covers this range of values.

Hotspot geometry. In the ST-U configuration, the NS features two independent hotspots (see fig. 1 of S. Vinciguerra et al. 2023). The priors related to hotspot position and size are kept broad: the prior distribution of ϕ covers a full cycle: $U(-0.25, 0.75)$, the θ is flat in cosine and goes from north pole to south pole: $\cos \theta \sim U(-1, 1)$, and the hotspot is allowed to have ζ of up to half the star: $U(0, 90)^{\circ}$. However, the two hotspots are mutually constrained in their positioning and sizes in that they do not overlap and they must be ordered, meaning that the primary hotspot has a smaller colatitude than the

secondary. In practice, all samples that do not adhere to these requirements are rejected before costly computation. In the ST-S configuration, the hotspots are identical and antipodal. That means that the number of hotspot parameters is halved because all parameters of the secondary spot are tied to the primary. In addition, $\cos \theta_p$ (the primary hotspot) is constrained to the upper hemisphere $U(0, 1)$, given that the NS is north-south symmetric.

Similar to B. Dorsman et al. (2025, 2026), we do not adopt any prior information on the hotspot spectral parameters T_{seed} , T_e , and τ . Instead, we sample uniformly from the full table from A. Bobrikova et al. (2023). For inner disc temperature (T_{in}), we use the same broad prior as B. Dorsman et al. (2026) from 0.01 to 0.6 keV. However, the shape is slightly changed because here we use a distribution flat in keV, whereas previously it was flat in $\log_{10}(\text{K})$. For R_{in} , we also use the same prior as B. Dorsman et al. (2026): a flat distribution between R_{eq} and the corotation radius $R_{\text{co}} = (GM/4f^2\pi^2)^{1/3}$, where G is the gravitational constant. As a result, the prior for R_{in} depends on other model parameters: R_{eq} , M , and the constant f . In practice, during sampling the value for R_{in} is drawn from within a 5–40 km range, and samples that are outside the range $R_{\text{eq}}-R_{\text{co}}$ are rejected before any costly computation. Other works have found a breadth of values for R_{in} . S. V. Molkov et al. (2024) obtained 24.6 km based on fitting of the pulse profile that includes eclipses by the disc, but they did not fully explore the parameter space to obtain errors. $11 R_g$ ($22.8 M_{1.4}$ km, where $M_{1.4} = 1.4 M_{\odot}$) is found by M. Mandal et al. (2025), and $10.0^{+1.2}_{-1.7} R_g$ ($21^{+3}_{-4} M_{1.4}$ km) or $6 R_g$ ($12 M_{1.4}$ km) by C. Malacaria et al. (2025), where the latter is obtained if fixing $i = 53^{\circ}$. We note however that R_{in} is expected to be variable throughout the evolution of the outburst, and instruments have taken measurements at different times. We opt to stick with the broad prior described above.

Finally, we include α -factors, scaling factors that are multiplied with the instrument responses to obtain the effective instrument responses, to account for potential deviation in cross-calibration. This is first done in the independent analysis of *IXPE* data and combination method 2 (whenever *IXPE* data are analysed) because *IXPE* has three instruments on board. Here, the prior for α_{DU1} , α_{DU2} , and α_{DU3} is $U(0.8, 1.2)$. Later, in combination method 3, where *IXPE* and NICER are jointly analysed, α are used again but are modified. There, we introduce α_N with prior $U(0.8, 1.2)$, but we avoid increasing the parameter space by fixing α_{DU1} to 1. While arbitrary, this restriction is a reasonable trade-off for computational efficiency, because it avoids an extra parameter that scales the flux, besides D . To marginally shrink parameter space further, we also shrink the α_{DU2} and α_{DU3} priors to $U(0.95, 1.05)$. This additional constraint is justified, because the *IXPE*-only analysis resulted in highly correlated α -factors (Fig. B2), meaning $\frac{\Delta \alpha_{\text{DU2}}}{\Delta \alpha_{\text{DU1}}} \approx \frac{\Delta \alpha_{\text{DU3}}}{\Delta \alpha_{\text{DU1}}} \approx 1.0$.

3 DATA PREPARATION

This section describes the preparation of the NICER and *IXPE* data into pulse profiles. Fig. 1 shows the outburst as observed by the Monitor of All-sky X-ray Image (MAXI, M.atsuoka et al. 2009), NICER, and *IXPE*. Only the subset of NICER data selected for analysis is shown. Figs 2 and 3 show, respectively, the pulse profiles and polarized pulse profiles resulting from the data preparation process.




Figure 1. The outburst as captured by *MAXI*, *IXPE*, and *NICER*. Only the *NICER* data selected for the analysis are shown.

3.1 NICER observations

Here, we outline the data processing steps undertaken for the *NICER* observations used in the PPM analysis of SRGA J144459.2–604207. These are largely similar to those presented in B. Dorsman et al. (2026), which analysed SAX J1808.4–3658. *NICER* is well suited for the PPM analysis given its large effective area ($\sim 1900 \text{ cm}^2$ at 1 keV) and exquisite time-tagging capabilities in the 0.2–12.0 keV energy range (100 ns absolute timing accuracy, K. C. Gendreau et al. 2016; G. Prigozhin et al. 2016). *NICER* is an external payload on the International Space Station comprising 56 co-aligned X-ray concentrator optics and silicon drift detectors in focal plane modules (FPMs), of which 52 are operational (B. LaMarr et al. 2016).

For this work, we utilized the *NICER* observations taken between 2024 February 21 to 2024 March 14, with OBSIDs 6 204 190 101 to 6 204 190 103 and 6 639 080 101 to 6 639 080 113 inclusive. However, we note that the observations in OBSIDs 6 639 080 105 to 6 639 080 112 suffered from large undershoot rates of over $700 \text{ c s}^{-1} \text{ FPM}^{-1}$, most likely due to high levels of optical loading,³ and thus did not provide any scientifically useful data (i.e. filtered exposure time of 0 s). The remaining data have very low undershoot rates of $\sim 3 \text{ c s}^{-1} \text{ FPM}^{-1}$.

We processed the data with the *NICER* Data Analysis Software (*NICERDAS*) v13 in *HEASOFT* v6.34 with the *NICER* calibration database (*CALDB*) version XT20240206. To generate the good time intervals for the PPM analysis, we adopted the following filtering criteria: Earth limb elevation angle of $\text{ELV} > 20^\circ$; bright Earth limb angle of $\text{BR_EARTH} > 30^\circ$; angular pointing offset (to the source) of $\text{ANG_DIST} < 54 \text{ arcsec}$; an undershoot rate per FPM of $0\text{--}500 \text{ c s}^{-1}$; an overshoot rate per FPM of $0\text{--}30 \text{ c s}^{-1}$, and a magnetic cut-off rigidity of $> 1.5 \text{ GeV c}^{-1}$. Additionally, we visually identified and manually excised five time intervals (a total of 349 s) that contained type-I X-ray bursts (for more on the bursts, see S. V. Molkov et al. 2024, T. Fu et al. 2025; A. Papitto et al. 2025). Finally, we have 24.8 ks of filtered exposure for our analysis.

To generate the pulse profiles for our PPM efforts, we first constructed a local timing solution by starting from the solution given

in A. Papitto et al. (2025) and utilized *PINT* (J. Luo et al. 2021) to model any remaining timing residuals. We derived a timing solution that is consistent (to within $1\sigma\text{--}2\sigma$) of the published timing solution (A. Papitto et al. 2025). The selected *NICER* data are recorded during the peak and early decrease of the outburst, as shown in Fig. 1. During this time, the pulse phase is stable, consistent with $\dot{f} = 0$, similar to the finding of A. Papitto et al. (2025). For the construction of the pulse profile, any evolution of flux and pulse amplitude are averaged out. However, we note that the amplitude of the first harmonic (the fundamental frequency) appears to rise slightly from 2.4 ± 0.4 per cent to 3.3 ± 0.4 per cent, while the flux decreases from ~ 350 to 250 c s^{-1} .

Similarly to B. Dorsman et al. (2026), we assessed whether constructing an averaged representative ancillary response file (ARF) or redistribution matrix file (RMF) would suffice for our work, or whether having observation-specific ARFs/RMFs would be more appropriate. We constructed the ARFs/RMFs for the individual observations by utilizing the *NICERARF* and *NICERRMF* tools⁴ (detectors 14 and 34 disabled). For the ARFs (i.e. effective area curves), the largest fractional difference between the observations used in this work was roughly 0.2 per cent. For the RMFs, there was virtually no difference⁵ between the different observations. Thus, it is appropriate to sum the ARFs and RMFs together into an averaged ARF/RMF for our analysis. The operation was facilitated by the *HEASOFT* tasks, *ftaddrmf* and *ftaddarf*, where we constructed the exposure time-weighted ARF and RMF.

We inspected the environmental and instrumental background spectrum, estimated by the 3C50 model (R. A. Remillard et al. 2022), in the energy range 0.3–10 keV and found that it varied from 0.5 per cent to 4.5 per cent throughout the observation. This fraction is acceptably small, and these estimated background counts are included in the pulse profile data but not separately accounted for during PPM, similarly to B. Dorsman et al. (2026). Finally, very few counts, ~ 1 per cent, are detected below channel 100 ($\sim 1 \text{ keV}$). To save on computational effort calculating low-energy radiation, we cut out the low-energy channels below channel 100.

3.2 IXPE observations

For this work, we utilize also the *IXPE* data, extracted from the same observation period as in A. Papitto et al. (2025), i.e. starting shortly after the *NICER* observation finished (*IXPE* GO 03250101; PI: A. Papitto). *IXPE* is well suited for spectropolarimetric timing studies owing to its imaging X-ray polarimetry capabilities in the 2–8 keV energy band, combining three co-aligned mirror modules with a net effective area of $\sim 100 \text{ cm}^2$ and event-by-event measurements providing spatial, spectral, timing, and polarization information, with $\lesssim 10 \mu\text{s}$ timing resolution (P. Soffitta et al. 2021; M. C. Weisskopf et al. 2022).

In contrast to the approach chosen for the *NICER* data, we instead adopted here the same timing solution as A. Papitto et al. (2025) to construct the pulse profiles. The observation lasted $\sim 10 \text{ d}$ and we used the *xpbin* tool in *IXPEOBSSIM* package version 30.6.4 (L. Baldini et al. 2022) to extract weighted *I*, *Q*, and *U* spectra (A. Di Marco et al. 2022) along with the *IXPE* *CALDB* v013 of 2024 January (F. Kislat et al. 2015). Source photons are

³https://heasarc.gsfc.nasa.gov/docs/nicer/analysis_threads/undershoot-intro/

⁴https://heasarc.gsfc.nasa.gov/docs/nicer/analysis_threads/arf-rmf/

⁵50 bins out of 6 million bins in element number-energy space showed discrepancies.




Figure 2. X-ray count data of NICER (left) and the three *IXPE* detector units (right). Top panels show the bolometric counts. Bottom panels show the phase-energy resolved pulse profiles. For *IXPE*, only the phase-energy resolved pulse profile of DU₁ is shown, because the other two are very similar.




Figure 3. Stokes *U* (left) and *Q* (right) data recorded by the three *IXPE* detector units. Top panels show the normalized (i.e. divided by counts in each phase-energy bin) bolometric Stokes parameters with 1σ error bars. Bottom panels show the phase-energy resolved Stokes parameters. For simplicity only, the polarized pulse profile of DU₁ is shown.

extracted from a circular region with a radius of 100 arcsec centred at the source position; considering the source flux, background rejection or subtraction is not needed (A. Di Marco et al. 2023). The model-independent *pcube* algorithm was employed to determine the polarization quantities as a function of pulse phase in A. Papitto et al. (2025). Here, we instead performed a spectropolarimetric analysis in 16 phase bins. For the PPM analysis, we selected only phase-energy bins with nominal energies in the 2–6 keV range, where the number of *I* counts in each bin is $\gtrsim 20$ without additional rebinning. In this regime, the Stokes parameters *Q* and *U* are approximately normally distributed, and

we therefore model their likelihoods with Gaussian functions using the measured uncertainties as standard deviations (see L. Baldini et al. 2022, for details).⁶ Fig. 3 shows that the polarization signature of the data is small but significant, which is consistent with theoretical predictions, such as by A. Bobrikova

⁶For fewer counts per bin, a Gaussian function is not a good approximation for the distribution of counts. A Poisson function is also not valid due to the presence of negative counts. However, Gaussian functions could still be used after rebinning or excluding nearly empty bins.

et al. (2023). The phase dependence is not significant however, and the polarization signature is consistent with being phase-independent. A. Papitto et al. (2025) reached the same conclusions in their analysis of the data. We note that the polarization here is smaller than in Figs 7 and 8 of A. Papitto et al. (2025), which results from the use of different energy bands (3–6 keV) and analysis tools. The normalized q , u values plotted in A. Papitto et al. (2025) did not include acceptance corrections. However, their actual model fitting used unnormalized Q , U and accounted for the instrumental response. Those fit products are therefore more similar to the Stokes products used in our analysis than the normalized p -cube curves shown in the A. Papitto et al. (2025) figures.

4 RESULTS

This section presents the results from the PPM analyses of the NICER and *IXPE* data. First, the data sets are analysed individually in Sections 4.1 and 4.2. Next, we combine analysis of both data sets in multiple ways. In Section 4.3, we combine the likelihood surfaces of the separate analyses with combination method 1, as described in Section 2. Following this, in Section 4.4, we apply combination method 2. Table 2 gives an overview of the inferred parameter values that result from these methods, which are the main results of this work. However, in Appendix A, we explored also the usage of combination method 3. The results of that are not included in the main body of the paper because only the less preferred ST-S model could be used. An overview of those results is given in Table A1. These two aforementioned tables specify also relevant sampling parameters: number of parameters, the nested sampling efficiency, run cost, and number of live points. The text related to each run discusses the choice in live points and convergence.

4.1 NICER-only analysis

For this analysis 2000 live points and an X-PSI sampling efficiency of 0.1 were used. Another run of 5000 live points was attempted and intermediate results were converging to the same parameters as were found here. We deem the current settings are thus sufficient and that the 2000-live-point run has converged. The 5000-live-point run was not completed due to expected high computational cost.

As shown in Fig. 4, analysis of the NICER data leads to a visually good fit, apart from spectral residuals at ~ 1.7 keV that are likely due to unmodelled NICER spectral edges: Aluminium K edge/fluorescence (from the detector window) at 1.56 keV and Silicon K edge (from the window and bulk detector) at 1.84 keV.⁷ Because of these NICER spectral edges, formally the fit is very poor, with $\chi^2/\text{DOF} = 30108/28781 = 1.0461$, corresponding to a p -value of 2.4792×10^{-8} . The p -value quantifies the probability of obtaining a data set with a χ^2 value at least this large, given the model with these parameters ($\chi^2 = \sum_i (m_i - d_i)^2/m_i$, where d_i and m_i are data and model counts in each phase-energy resolved bin i – more specifically, m_i are the posterior – expected model counts based on 100 draws from the posterior).

Three of the parameters ($M = 2.17^{+0.02}_{-0.03} M_\odot$, $R_{\text{eq}} = 13.82^{+0.10}_{-0.14}$ km, and $\zeta_p = 88.6^{+0.9}_{-1.2}$) tend towards the maximum of the prior, whereas others ($D = 2.88^{+0.05}_{-0.04}$ kpc, $i = 50.9^{+0.8}_{-0.5}$, and $\theta_p = 1.1^{+0.3}_{-0.2}$) tend towards the minimum of the prior (see Fig. 5).

It is notable that many parameters peg to the bounds of their respective priors, which could be a hint that parameter space outside the current prior distributions are worth exploring. We discuss this possibility in Section 5.3.3. The corresponding spot geometry has a large, bright primary hotspot that produces a significant number of non-pulsed counts, while wobbling slightly around the rotational axis, as shown by the projection in Fig. 4. This geometry is very similar to what was found in B. Dorsman et al. (2025) for SAX J1808.4–3658. R_{in} is also very large, $32.6^{+0.3}_{-0.5}$ km and near R_{co} , the upper bound of the prior. In combination with the inferred observer inclination (i), this implies that no rays from the secondary hotspot are blocked from the observer.

The spectrum in Fig. 4 shows that the accretion disc dominates at low energy but becomes subdominant at 1.3 keV. Above ~ 2 keV, the primary hotspot is responsible for over 90 per cent of the counts. The secondary hotspot, despite the fact that it creates considerably fewer counts than the primary, is almost entirely responsible for shape of the pulsation.

4.2 IXPE-only analysis

The *IXPE* data consisted of Stokes I , Q , and U data. For comparison, we performed *IXPE*-only analyses both with and without fitting the Stokes I data. Unlike the NICER analysis, for the final *IXPE* analyses the accretion disc model was not applied. We did a test including the accretion disc for the Stokes IQU -analysis, but the run was much slower and resulted in a worse evidence than a run without the disc with the same sampler settings (1000 live points and 0.3 sampling efficiency). Presumably good fits are found without the disc because its contribution to the *IXPE* counts is negligible given *IXPE*'s harder energy range compared to NICER.

For the *IQU* results presented in Table 2 and Fig. 5, we used 4000 live points and 0.1 sampling efficiency. With the initial sampler settings, the mode corresponding to a smaller R_{eq} and higher M was not recovered (see Fig. 5). The mode appeared when increasing the number of live points from 1000 to 4000, whereas reducing the sampling efficiency from 0.3 to 0.1 did not lead to significant further shifts in the posteriors. Nevertheless, slight broadening of many CIs was still observed, and the convergence was thus not formally proven. That would have been computationally expensive to do, because the last run spent already almost 50 kilo core-hours (kch). For the Stokes Q - and U -only analysis, further improved sampler settings were used (10 000 live points and 0.1 sampling efficiency), due to the very short run time (less than 1 kch) of the much noisier data. The full posterior distributions from these runs are shown in Appendix B.

The QU and IQU analyses result in visually good fits in all the fitted phase-resolved Stokes spectra (see e.g. Stokes I for DU₁ in Fig. 4). For the Stokes I fit statistics (combined over detector units), we get $\chi^2/\text{DOF} = 4940/4827$, corresponding to a p -value of 0.12632. As in A. Papitto et al. (2025), the *IQU* fit is dominated by the Stokes I signal since the polarization degree is rather low (< 5 per cent) and the polarization angle is roughly constant in phase.

The results generally agree with those obtained in A. Papitto et al. (2025). In particular, the inferred geometry matches that of A. Papitto et al. (2025) when we analyse only the Q and U data, although with large CIs. In this case, we find, for example, that $i = 77^{+8}_{-10}$, $\theta_p = 17^{+50}_{-9}$, and $\theta_s = 156^{+17}_{-50}$. Slightly different

⁷https://heasarc.gsfc.nasa.gov/docs/nicer/analysis_threads/arf-rmf/

Table 2. An overview of the analyses done in this work and the inferred parameters. In the header, various analysis and fit statistics are shown. The sampling efficiency is the nested sampling efficiency as defined in X-PSI. This number is modified before being passed to MULTINEST, because MULTINEST is unaware some samples are rejected in the prior. χ^2 is calculated using 100 draws from the posterior, as described in the text. The $\hat{\chi}^2 = \chi^2/\text{DOF}$, and the p -value are given based on the X-ray counts without taking into account polarization. For the *IXPEI*, *Q*, and *U* runs, the χ^2/DOF and the p -value are given for the combined Stokes *I* data of the DU₁, DU₂, and DU₃ detectors. For the KDE combination, the sampling parameters of the last sampling step are given. The body of the table gives an overview of the posterior median values and 68 per cent CIs of all parameters.

Name	NICER-only	<i>IXPEI</i> -only (<i>IQU</i>)	<i>IXPEI</i> -only (<i>QU</i>)	Combination method 1	Combination method 2		
Data	NICER	<i>IXPEIQU</i>	<i>QU</i>	NICER, <i>IXPEIQU</i>	NICER	NICER	<i>IXPEIQU</i>
Prior updated by	–	–	–	–	<i>IXPEIQU</i>	<i>IXPEIQU</i> , flat N_H	NICER
Hotspot config.	ST-U	ST-U	ST-U	–	ST-U	ST-U	ST-U
Sampling aspects							
Num. params.	19	21	21	5	19	19	21
Samp. eff. (X-PSI)	0.1	0.1	0.1	0.1	0.1	0.1	0.3
Run cost (kch)	2.9	45.8	0.4	$\lesssim 0.1$	9.3	18.5	13.3
Live points	2000	4000	10 000	4000	2000	2000	4000
$\ln(\mathcal{Z})$	−103760	−55372	−43479	−22.792	−103800	−103740	−55369
Fit quality							
χ^2/DOF	30108/28781	4940/4827	–	–	30204/28781	30103/28781	4953/4827
$\hat{\chi}^2$	1.0461	1.0233	–	–	1.0494	1.0459	1.0262
p -value	2.4792×10^{-8}	1.2632×10^{-1}	–	–	2.6326×10^{-9}	2.8227×10^{-8}	9.9836×10^{-2}
M [$M_\odot$]	$2.17^{+0.02}_{-0.03}$	$2.02^{+0.12}_{-0.36}$	$1.4^{+0.4}_{-0.3}$	$1.928^{+0.012}_{-0.012}$	$2.15^{+0.03}_{-0.06}$	$2.15^{+0.03}_{-0.05}$	$2.16^{+0.02}_{-0.03}$
R_{eq} [km]	$13.82^{+0.10}_{-0.14}$	$10.7^{+2.4}_{-1.0}$	$13.1^{+0.6}_{-0.8}$	$13.94^{+0.04}_{-0.05}$	$11.1^{+0.3}_{-0.3}$	$13.6^{+0.3}_{-0.3}$	$13.85^{+0.09}_{-0.11}$
D [kpc]	$2.88^{+0.05}_{-0.04}$	$3.3^{+0.3}_{-0.3}$	$3.3^{+0.8}_{-0.5}$	$3.060^{+0.017}_{-0.019}$	$2.71^{+0.07}_{-0.07}$	$2.55^{+0.05}_{-0.03}$	$2.88^{+0.04}_{-0.03}$
i [deg]	$50.9^{+0.8}_{-0.5}$	63^{+12}_{-8}	77^{+8}_{-10}	$54.2^{+0.2}_{-0.3}$	53^{+3}_{-2}	$51.1^{+1.4}_{-0.8}$	$50.8^{+0.7}_{-0.5}$
N_H [10^{21}cm^{-2}]	$27.82^{+0.07}_{-0.08}$	$22.1^{+1.0}_{-0.9}$	24^{+3}_{-3}	$27.59^{+0.03}_{-0.03}$	$26.340^{+0.008}_{-0.015}$	$28.04^{+0.10}_{-0.11}$	$27.82^{+0.07}_{-0.08}$
ϕ_p [cycles]	$-0.34^{+0.05}_{-0.04}$	$-0.06^{+0.14}_{-0.17}$	$0.28^{+0.16}_{-0.70}$	–	$-0.45^{+0.02}_{-0.02}$	$-0.32^{+0.27}_{-0.07}$	$-0.25^{+0.05}_{-0.03}$
θ_p [deg]	$1.1^{+0.3}_{-0.2}$	$0.7^{+0.4}_{-0.3}$	17^{+50}_{-9}	–	$3.0^{+1.4}_{-0.9}$	$1.1^{+0.3}_{-0.3}$	$1.5^{+0.4}_{-0.3}$
ζ_p [deg]	$88.6^{+0.9}_{-1.2}$	82^{+5}_{-10}	43^{+15}_{-29}	–	86^{+3}_{-3}	86^{+3}_{-3}	83^{+5}_{-7}
$T_{\text{seed},p}$ [keV]	$0.822^{+0.009}_{-0.006}$	$0.85^{+0.05}_{-0.14}$	$1.39^{+0.11}_{-0.60}$	–	$0.97^{+0.03}_{-0.03}$	$0.795^{+0.014}_{-0.013}$	$0.649^{+0.008}_{-0.009}$
$T_{e,p}$ [keV]	$25.9^{+0.6}_{-0.5}$	76^{+15}_{-13}	60^{+20}_{-20}	–	$23.0^{+0.9}_{-0.8}$	$24.2^{+0.8}_{-0.8}$	40^{+4}_{-3}
τ_p [–]	$2.003^{+0.019}_{-0.023}$	$1.18^{+0.06}_{-0.07}$	$2.0^{+0.8}_{-0.7}$	–	$2.15^{+0.04}_{-0.04}$	$2.09^{+0.03}_{-0.03}$	$1.39^{+0.05}_{-0.05}$
ϕ_s [cycles]	$0.080^{+0.006}_{-0.006}$	$-0.311^{+0.049}_{-0.012}$	$-0.1^{+0.3}_{-0.2}$	–	$0.034^{+0.016}_{-0.014}$	$0.082^{+0.016}_{-0.009}$	$-0.286^{+0.010}_{-0.012}$
θ_s [deg]	145^{+3}_{-3}	146^{+13}_{-10}	156^{+17}_{-50}	–	143^{+5}_{-5}	146^{+5}_{-4}	139^{+8}_{-13}
ζ_s [deg]	52^{+3}_{-3}	22^{+8}_{-5}	51^{+18}_{-36}	–	59^{+7}_{-6}	51^{+7}_{-28}	13^{+9}_{-4}
$T_{\text{seed},s}$ [keV]	$0.84^{+0.04}_{-0.04}$	$1.34^{+0.12}_{-0.16}$	$1.3^{+0.2}_{-0.5}$	–	$0.99^{+0.06}_{-0.06}$	$0.83^{+0.58}_{-0.08}$	$1.29^{+0.17}_{-0.26}$
$T_{e,s}$ [keV]	61^{+9}_{-9}	86^{+11}_{-23}	60^{+20}_{-20}	–	47^{+9}_{-8}	75^{+16}_{-16}	76^{+18}_{-25}
τ_s [–]	$3.29^{+0.12}_{-0.18}$	$2.8^{+0.5}_{-0.6}$	$1.9^{+0.8}_{-0.7}$	–	$3.25^{+0.16}_{-0.30}$	$3.22^{+0.19}_{-0.27}$	$2.6^{+0.6}_{-0.8}$
T_{in} [keV]	$0.352^{+0.004}_{-0.003}$	–	–	–	$0.413^{+0.008}_{-0.007}$	$0.336^{+0.004}_{-0.003}$	–
R_{in} [km]	$32.6^{+0.3}_{-0.5}$	–	–	–	$19.8^{+0.8}_{-0.7}$	$32.5^{+0.4}_{-0.7}$	–
α_{DU1} [–]	–	$1.10^{+0.05}_{-0.08}$	$1.04^{+0.10}_{-0.12}$	–	–	–	$0.99^{+0.11}_{-0.08}$
α_{DU2} [–]	–	$1.13^{+0.05}_{-0.08}$	$1.02^{+0.11}_{-0.12}$	–	–	–	$1.01^{+0.11}_{-0.09}$
α_{DU3} [–]	–	$1.10^{+0.05}_{-0.08}$	$1.02^{+0.11}_{-0.12}$	–	–	–	$0.99^{+0.11}_{-0.08}$
χ_0 [rad]	–	$-0.44^{+0.18}_{-0.21}$	$-0.46^{+0.10}_{-0.09}$	–	–	–	$-0.6^{+0.6}_{-0.5}$

constraints are obtained when analysing all *I*, *Q*, and *U* data, with bimodal posterior distributions in many parameters.⁸ In this case, we find $i = 63^{+120}_{-8}$, $\theta_p = 0.7^{+0.40}_{-0.3}$, and $\theta_s = 146^{+130}_{-10}$, and notably a larger primary hotspot size, with ζ_p covering almost half of the

NS surface (similar to the NICER analysis). In both cases the NS mass and radius are only weakly constrained.

4.3 Combination method 1: combining independent posteriors

This section presents the results of combining the posteriors of the *IXPEI*-only (Stokes *I*, *Q*, and *U*) with the NICER-only

⁸Preliminary tests fitting only Stokes *I* yielded results very similar to those obtained when fitting all Stokes parameters.




Figure 4. Projection plots, spectra, pulse profiles, and residuals for the (from left to right) NICER-only, *IXPE*-informed NICER, *IXPE*-only, NICER-informed *IXPE* analyses. The top row of panels show the projection plots of the maximum a posteriori (MAP) samples from the perspective of the observer. The solid lines indicate the primary hotspots, the dash-dotted lines indicate the secondary hotspots, and the dashed lines indicate inner disc radii. The crosses/dots indicate the centres of the hotspots, which are on the observer/opposite side of the star, respectively. The accretion disc is not present in the *IXPE* models. Relativistic effects are not shown in these projection plots. Stars indicate the Northern and Southern rotational poles. The second row of panels show the observed phase-summed spectra of the MAP samples. Here, the black lines indicate data spectra, and the grey dashed lines indicate the posterior-expected model spectra. The solid, dashed and dash-dotted lines are the same as in the top panels and sum up to the total model spectrum in black. The third row of panels show the intrinsic source spectra before circumstellar and interstellar absorption and before convolving with the instrument response. The lines correspond to the same component as the above panels. The fourth row of panels show the bolometric pulse profile of the data and posterior-expected model in black and grey, respectively. The model also shows the 16 and 84 per cent Poisson error CIs in light grey. The upper filled-in region displays the bolometric contribution of the secondary hotspot, and the lower region the contribution of the primary hot spot (and disc, if present). Finally, the bottom row of panels show the phase-energy resolved residuals, where m is the model counts and d is data counts in each phase-energy bin. In the case of the *IXPE* analyses, only the spectra, pulse profiles, and residuals of DU₁ are shown (because DU₂ and DU₃ look very similar), although all three were fitted simultaneously during sampling.

analysis. We used KDEs to add the ln-likelihood surfaces, following the methodology outlined in Section 2.2. As a reminder, here only the constant parameters (θ_c) are combined, because all potentially variable parameters may change between the recording of the two data sets. For the final sampling step, we used a sampling efficiency of 0.1 and 4000 live points.

Fig. 5 shows a corner plot of the inferred posteriors of θ_c . Interestingly, the combined run results in a lower $M = 1.928^{+0.012}_{-0.012}$

$M_\odot$ and higher $R_{\text{eq}} = 13.94^{+0.04}_{-0.05}$ km than either of the individual runs. At first glance, this result seems counter-intuitive, but we interpret it by looking at the $M - R_{\text{eq}}$ posteriors. Here, we see that the NICER-only posterior lies in the top-right corner, whereas the *IXPE*-only posterior disfavors this corner in favour of either lower M or lower R_{eq} . It appears that the best overlap found during sampling was towards lower M and thus the combined posterior is found there. The CIs of the combination are noticeably




Figure 5. Posterior distributions of the constant parameters from the NICER- and *IXPE*-only analysis, and their KDE combination (method 1, see Section 4.3). The dashed and solid lines on the diagonal show the 1D marginal priors and posteriors of each parameter, respectively. For visibility, these are scaled such that the distributions peak at the same height. The shaded intervals correspond to the 68.3 per cent CIs also listed above the diagonal. The off-diagonal plots show 2D marginal posteriors for each parameter combination. There, contours correspond to the 68.3 per cent, 95.4 per cent, and 99.7 per cent credible regions.

narrower than posterior of either constituent runs, as is expected from the multiplication of two peaked distributions, and more so if they are partially disjoint.

Based on the change in N_{H} , we consider the possibility that N_{H} is not a constant parameter that, instead, has decreased between the NICER to the *IXPE* observations. To cover for that possibility, we exclude the N_{H} from the combination and rerun the final sampling step in the combination process. This moves the posteriors significantly (beyond their 68 per cent CIs), with now $M = 1.885^{+0.020}_{-0.015} M_\odot$ and $R_{\text{eq}} = 13.84^{+0.05}_{-0.06}$ km, although this change is small compared to the width of both the NICER- and *IXPE*-only posteriors.

4.4 Combination method 2: two-stage Bayesian analysis

This section describes the two-stage Bayesian analyses, where the posteriors of the constant parameters that result from the analysis of one data set are used as priors in the analysis of the other data set. This analysis is done in two ways: using *IXPE*-only (*IQU*) posteriors as priors for subsequent NICER analysis, and vice versa.

4.4.1 *IXPE*-informed NICER analysis

In this analysis, the posteriors that resulted from the *IXPE*-only analysis were implemented as updates to the priors of the con-

stant parameters for the NICER analysis. This analysis was done with 2000 live points and a sampling efficiency of 0.1, like the NICER-only analysis. Compared to 1000-live-point runs, these results did not change much, indicating that the sampler settings were sufficient for these runs to be converged.

As shown in Table 2, M is consistent with the NICER-only result at $2.15^{+0.03}_{-0.06} M_\odot$, but R_{eq} shrinks significantly to $11.1^{+0.3}_{-0.3}$ km and is consistent with the *IXPE*-only result. R_{in} shrinks to $19.8^{+0.8}_{-0.7}$ km, which implies a brighter disc and the partial occlusion of the emission of the secondary hotspot. The second column in Fig. 4 shows that the hotspot configuration closely resembles the configuration found in the NICER-only analysis. The secondary hotspot is brighter compared to NICER-only, owing to its higher T_{seed} and ζ . Unlike before, we can now clearly see that the reduction in flux between NICER and *IXPE* observations (see Fig. 4) is mainly due to the decrease in $T_{\text{seed,p}}$. This is consistent with the idea that the accretion rate is higher during the NICER observation.

$\mathcal{Z}$ is significantly worse for this model, with a penalty of 41.19 in $\ln$ -space (see e.g. R. E. Kass & A. E. Raftery 1995), and the fit is also significantly worse, with the χ^2/DOF increasing to $30204/28781 = 1.0494$ and the p -value reducing by around a factor of 10 to 2.6326×10^{-9} . These metrics indicate that this solution is significantly disfavoured by the NICER data compared to the NICER-only solution.

To allow for the possibility that N_{H} is not a constant parameter (but evolves during the outburst), we also ran the *IXPE*-informed NICER analysis while keeping the N_{H} prior from Table 1. The solution found in this case is better than the previous *IXPE*-informed NICER result, and marginally better than even the NICER-only result, as shown by $\ln(\mathcal{Z})$ and the χ^2 value in Table 2. This finding supports the idea that N_{H} might not be constant, but instead may have decreased between the time of the NICER observations and the *IXPE* observations. Compared to the previous *IXPE*-informed NICER result, M is effectively unchanged, but R_{eq} increases to $13.6^{+0.3}_{-0.3}$ km and is now more consistent with the NICER-only result. This improved consistency is also true for the other parameters, such as N_{H} and the hotspot configuration. This experiment shows that a good prior on N_{H} is important for the inference of R_{eq} .

4.4.2 NICER-informed *IXPE* analysis

In this analysis, the posteriors that resulted from the NICER analysis were implemented as updates to the priors of the constant parameters for the *IXPE* (*IQU*) analysis. While the number of live points is still 4000, we note though that for computational speed-up, the sampling efficiency was set in this case to 0.3 instead of 0.1.

Interestingly, the $\ln(\mathcal{Z})$ found here is 3 higher in $\ln$ -space than in the *IXPE*-only analysis, which is notable but not definitive evidence in favour of this model. However, the χ^2/DOF fit is marginally worse at $4953/4827 = 1.0262$. As shown in Table 2, the constant parameters are mostly consistent with the updated prior (the NICER-only posterior). Here, M is inferred to be $2.16^{+0.02}_{-0.03} M_\odot$ and R_{eq} is inferred to be $13.85^{+0.09}_{-0.11}$ km. Similarly as for the *IXPE*-informed NICER analysis, we see that the reduction in flux compared to the NICER observation is mainly caused by a reduction in $T_{\text{seed,p}}$.

The variable parameters are mostly consistent with the *IXPE*-only posteriors. Fig. 6 shows that, as we had speculated in Sec-




Figure 6. Posterior distributions for the *IXPE* Stokes *I*, *Q*, and *U* analysis informed by NICER posteriors (IQU_prNICER), and two NICER analyses informed by *IXPE* posteriors. The first retains the flat N_H prior (NICER_prIQU_flat_nH) and the second utilizes the *IXPE* posterior of N_H as prior (NICER_prIQU). These correspond to combination method 2 (see Section 4.4). See the caption of Fig. 5 for a detailed description of the contents of the diagonal and off-diagonal panels.

tion 2.2, rather than leading to consistent results, there are differences between the results of the two-step analyses, depending on which data set is analysed first.

5 DISCUSSION

In this section, we will discuss the obtained results and this discussion is structured as follows. Section 5.1 starts out by providing an overview of obtained system parameters and compares them to the literature where possible. Multiple methods of data analyses were explored – in which the two data sets were considered both separately and together – so, Section 5.2 discusses reliability and fit quality of each analysis. Finally, Section 5.3 gives recommendations for future work.

5.1 Overview of inferred parameters and literature comparison

This section provides an overview of inferred parameters. These are grouped into parameters that govern hotspots, mass and radius, geometry, pulse profile, distance, the polarization, and the spectrum. We compare results against literature values where possible.

Hotspots. A consistent trend in the results is that the hotspots of SRGA J144459.2–604207 are very large: the primary hotspot occupying up to half of the stellar surface, and the secondary hotspot size in the range $\sim 20^\circ$ – 70° . We can compare this against J. J. E. Kajava et al. (2011) and T. Salmi et al. (2018), who found smaller hotspots for SAX J1808.4–3658. They used *RXTE* data that covers a higher energy range, however, where the effective

hotspots could look smaller if they have a temperature gradient. In addition, they accounted for disc reflection in their models, which we do not, and that might reduce their inferred hotspot size. B. Dorsman et al. (2026) also find generally large hotspots for SAX J1808.4–3658 using NICER data, although they were not as large when they used the STU-Disc configuration: $\zeta_p \approx 47^\circ$ – 59° . Simulations (e.g. M. M. Romanova et al. 2004; A. K. Kulkarni & M. M. Romanova 2013; P. Das et al. 2025) predict typically smaller hotspots, but perhaps more importantly they predict very different hotspot shapes than those allowed in our model space, e.g. crescents, rings, and elongated shapes. Although upgrades to hotspot shape modelling are needed, our results indicate that the ‘effective’ hotspots of AMPs could cover a significant portion (even up to half) of the stellar surface.

Mass and radius. The analyses presented here consistently yielded a high M at 1.9–2.2 $M_\odot$, where 2.2 $M_\odot$ is the upper bound of the prior. For R_{eq} , we observe that generally two solution modes are found. On the one hand, there is a high R_{eq} mode (13.4–14 km), associated with the NICER-only analysis, analyses with the NICER prior and the KDE combination. On the other hand, there is a low R_{eq} mode (9.7–13.1 km), preferred by the *IXPE*-only analysis and analysis with the *IXPE* prior (although Fig. 5 shows that the *IXPE*-only analysis also found a high radius mode). The inferred M and R_{eq} are near the prior boundaries in many cases, so we recommend exploration of the robustness of results against changes in M – R_{eq} priors. We discuss this point further in Section 5.3.3. We recommend this additional step should be done before including these present results in any new EoS inference.

Geometry. Taking our results together, we find the inclination of the system to be moderately high ($i = 50^\circ$ – 75°), the primary hotspot to be centred closely around the rotational axis ($\theta_p = 0^\circ$ – 4°), and the secondary centred on the opposing hemisphere ($\theta_s = 125^\circ$ – 159°). We note here that i at 50° is close to the lower boundary prior boundary, which was only based on inferred values in literature, so this motivates an investigation with a lower prior in future work. Furthermore, we found R_{in} to be distributed in the range of 19–33 km.

Let us compare these results against S. V. Molokov et al. (2024), who investigated the persistent pulsations discovered in the ART-XC data, which started just before and overlapped with the NICER observation. S. V. Molokov et al. (2024) fitted the pulse profile by eye, assuming two antipodal spots and fixed $M = 1.4 M_\odot$ and $R_{eq} = 12$ km. They used a simplified atmosphere model $I(\alpha) \propto 1 + h \cos \alpha$ where α is the angle relative to the local normal in the spot comoving frame and h is the anisotropy parameter $-1 < h < 0$. They found a good fit with $i = 58^\circ$, $\theta = 14^\circ$, $\zeta = 33^\circ$, and $R_{in} = 24.6$ km. While their i and R_{in} are in agreement with our results, their θ and ζ are not. This disagreement may be accounted for by their restriction to antipodal spots, simplified atmosphere model, and limited search of parameter space.

Let us also compare the results against A. Papitto et al. (2025), who performed simplified PPM on the *IXPE* Stokes *I*, *Q*, and *U* data. They assumed two antiphased point-like spots with an emission model consisting of power-law emission $I(E) \propto E^{-(\Gamma-1)}$, Γ fixed at 1.7, and same simplified anisotropy distribution with parameter h as S. V. Molokov et al. (2024). They also used a normalization factor to scale the flux to the data. Furthermore, they restricted their phase-resolved analysis to the 3–6 keV energy band, combining all energy channels within this interval to maximize the polarized signal, whereas we performed a fully phase- and energy-dependent analysis over the 2–6 keV range. Using fixed $M = 1.4 M_\odot$, $R = 10$ km, and $h = -0.3$, they found $i = 74.1^{+5.8}_{-6.3}$,

which is consistent with our *IXPE*-only Stokes I , Q , and U result, but inconsistent with any of the results where NICER data were involved (in which i is closer to 50°). Furthermore, they infer $\theta_p = 11.8^{+2.5}_{-3.5}$, and $\theta_s = 172.6^{+2.0}_{-1.0}$, which are somewhat in tension with our results, except the *IXPEQ* and U analysis (although all our inferred hotspots lie in the same general region on the NS surface). These differences in results can easily be accounted for by the various differences in modelling mentioned above.

Pulse profile and timing properties. In S. V. Molkov et al. (2024), the primary hotspot causes the main pulse and the secondary causes the interpulse, with a plateau due to partial eclipse by the accretion disc. NICER also observed the plateau (see Fig. 2), so it is unlikely that this feature can be attributed to Poisson noise. Motivated by S. V. Molkov et al. (2024) and unlike B. Dorsman et al. (2026), we have included occlusion of hotspot emission by the disc. However, our best fit to NICER data does not require disc blocking and still leaves some curvature in the pulse profile where the plateau should be. We judge our fit to be acceptable, because we see no clear phase-dependent structure in the residuals, despite that difference. Still, the *IXPE*-informed NICER fit (second column in Fig. 4) features a partially occulted secondary hotspot. In contrast, here the primary is in antiphase with the overall pulsation and is responsible for flattening the interpulse. In the end, we judge that the results are not conclusive and more investigation is warranted. Z. Li et al. (2025) also studied the pulsations and revised the spin evolution, favouring a spin-up episode during the NICER observation. They selected a longer time interval for this timing solution however, that we exclude due to high undershoot rates in the NICER data (see Section 3.1).

Distance. We found the system to be very close-by (2.5–3.3 kpc) compared to what has so far been inferred in the literature (~ 8 –11 kpc, see Section 2.3). We advise caution when using this value range, since we infer D near the lower prior bound. The lower bound was selected rather arbitrarily to cover the lower end of the distance estimate range of S. V. Molkov et al. (2024), so we do not associate our lower bound with a high level of confidence. A smaller distance corresponds to a brighter signal, so we speculate that the discrepancy with the literature is explained by us omitting a source of radiation in the model, such as the accretion column or disc reflection.

Absorption column. The analyses of the NICER and *IXPE* data inferred different values for N_H : $(22.1^{+1.0}_{-0.9}) \times 10^{21} \text{ cm}^{-2}$ for *IXPE*-only or $(26\text{--}28) \times 10^{21} \text{ cm}^{-2}$ if NICER-data are in any way included in the analysis. This may indicate that this parameter has evolved between these observations. This variation is an important consideration, because in the combined analyses, we found that treating this parameter either as constant or variable influences the inference of R_{eq} .

Polarization. A. Papitto et al. (2025) suggest that the most natural explanation for the polarization detected by *IXPE* is Comptonized emission from hotspots, with $T_e \sim 20$ –60 keV, $T_{\text{seed}} \sim 1$ keV, and moderate optical depth at $\tau \approx 1$ –2. This origin is mostly consistent with our findings in the analysis of *IXPE* data. We found T_e ranging from 37 to 97 keV, T_{seed} ranging from 0.65 to 1.5 keV, and τ from 1.1 to 1.4. The primary hotspot is typically towards the lower end of the range for these three parameters while the secondary is towards the higher end.

Spectrum. The broad-band X-ray spectrum was studied by C. Malacaria et al. (2025) with overlapping *XMM-Newton* and *NuSTAR* data. They found that an absorbed *bodyrad* ($T_{\text{bb}} = 0.64^{+0.02}_{-0.01}$ keV) and *diskbb* ($T_{\text{in}} = 0.44^{+0.09}_{-0.06}$ keV) and *nthComp*

($T_{\text{seed}} = 0.93^{+0.02}_{-0.02}$ keV) account well for the lower energy part ($\lesssim 4$ keV) of the spectrum. These observations took place just after the NICER observation, at the start of the *IXPE* observation. Our *IXPE* pulse profile is averaged over a longer part of the decline, so the results are not fully comparable. Still, we find some agreement. Their *diskbb* T_{in} is only a bit higher than our T_{in} (0.3 – 0.4 keV). Their *nthComp* T_{seed} is similar to our $T_{\text{bb,p}}$ and $T_{\text{bb,s}}$ during the NICER observation, although $T_{\text{bb,s}}$ increases in the later *IXPE* observation. Their T_e is a lower temperature ($18.2^{+0.7}_{-2.7}$ keV), which is close to our findings for $T_{e,p}$, but not for $T_{e,s}$, where our inferred value is significantly hotter. The differences could perhaps be explained by their more detailed modelling of reflection, which starts to become dominant at higher energy but is not accounted for in our model. Also notable is that they have another low-energy *bodyrad* component, which we do not have. Plausibly such a component could arise if the surface of the NS outside of the hotspots is hot enough.

For the higher energy ($\gtrsim 4$ keV) part of the spectrum, they found good fits by adding to the aforementioned lower energy components either a semiphenomenological model that features two *diskline* models or a comprehensive reflection model that accounts for the broadened iron line and other reflection from the disc. They infer different values between these models for i and R_{in} . Their model with two *diskline* components fits best at $i = 85^\circ$ and $R_{\text{in}} = 10.0^{+1.2}_{-1.7} R_g$ ($30^{+5}_{-4} M_2$ km, where $M_2 = 2 M_\odot$). However, their reflection model fits best with $i = 53^\circ$ and $R_{\text{in}} = 6 R_g$ ($18 M_2$ km), the latter being pegged at the lower limit. Our results are more in agreement with the smaller result for i , but both inferred values for R_{in} are consistent with our results, since we found both low and high values.

Z. Li et al. (2025) also analysed the X-ray broad-band spectrum using NICER, *Einstein Probe*, *IXPE*, *Insight-HXMT*, and *INTEGRAL* data. They found good spectral fits using absorbed *nthcomp* + *gaussian* in their spectral model to account for the NICER systematic at 1.7 keV. However, their $T_{\text{seed}} \sim 0.4$ keV and $T_e \sim 4$ –6 keV are much colder than our $T_{\text{seed}} \sim 0.8$ –1.1 keV and $T_e \sim 22$ –89 keV inferred for the NICER data. These diverging results are likely due to model choice, because we separately include a contribution from the disc blackbody at lower temperature. In their best-fitting model, *nthcomp* + *gaussian* + *relxillcp*, where the latter accounts for the broad-band accretion disc continuum and reflection, they still find similar T_{seed} and T_e as before.

5.2 Reliability and fit quality

Here, we discuss the reliability and, if applicable, fit quality of the results of each analysis. This discussion is subdivided as follows: Section 5.2.1 discusses the results of the stand-alone analysis of the NICER data, and Section 5.2.2 the same of the *IXPE*-data. Section 5.2.3 discusses the results of the combined data analyses.

5.2.1 PPM with NICER only

In this analysis, no clear pulsed structure remains in the residuals, which shows that the two circular hotspot model is sufficient to fit the data. However, the overall fit quality is not good, due to a non-pulsed structure remaining at ~ 1.7 keV, which we attribute to unmodelled NICER spectral edges.

We also noted that these data lead to tight parameter constraints for the PPM. However, several parameters are concentrated towards the bounds of their priors: M , R_{eq} , and ζ_p near

their respective upper bounds and D , i , and R_{in} near their lower bounds. The apparent tightness of these posterior distributions may be driven by prior bounds rather than the data. Accordingly, confidence in these results depends on the validity of the priors. We recommend reanalysis using alternative models and more conservative or better-justified priors to assess robustness.

5.2.2 PPM with IXPE only

In this case, good fits of the Stokes I data were found, indicating that the model is sufficient to fit the data, even without including an accretion disc. The parameter constraints are broader and generally overlapping with the NICER-only constraints, except for the posterior peaks of N_{H} and R_{eq} which are shifted to smaller values.

The polarization data were not constraining enough to affect the results much, compared to the constraints imposed by the IXPE Stokes I data. This finding is similar to that of A. Papitto et al. (2025), where most parameter information was also set by the total pulse profile alone. This may be explained by the relatively small polarization degree and very small phase-resolved variability of the polarization angle. Possibly, this will change for other stars and for future instruments like *eXTP*. The IXPE constraints on the geometry are also broader compared to the NICER-only constraints.

When comparing the analyses of the IXPE data versus the NICER data, it is clear that NICER data are more constraining. However, we consider the IXPE constraints important despite being broader than the NICER posteriors, because these evade potential biases specific to the analysis of NICER data, such as those related to systematics or the accretion disc (the latter was reported on by B. Dorsman et al. 2025). The Stokes Q and U results (see Figs B1 and B2) are a relevant point of comparison, meanwhile, if one is interested to avoid potential bias that may be caused by inaccurate modelling of the non-polarized emission.

5.2.3 PPM with combined NICER and IXPE data

We have seen that the two data sets give different answers on their own, even for parameters that should be the same, such as M . This discrepancy is intriguing and showcases that joint analysis is helpful to point out parameter evolution or inconsistency in inferred parameters. We explored two different methods to combine the data: (1) KDE-based combination of separate analyses, and (2) two-step Bayesian analyses. For discussion of the third method: a joint model, we refer to Appendix A, where this approach is explored to a limited degree.

Combination method 1 takes the samples of the two separate analyses and estimates the likelihood surfaces from them using KDE. A new sampling analysis is then done on the sum to estimate the final joint posterior. We found this method to be fast and affordable: only a few parameters, those assumed constant between observations, are used in the final sampling and the likelihood evaluation is fast. However, the results obtained with the KDE-based combination method may depend on the specific KDE hyperparameter choices used to reconstruct continuous likelihood surfaces from the two separate runs (NICER and IXPE), which are then combined to explore the shared parameter space. In this analysis, we adopted Scott's rule of thumb to set the KDE bandwidth (D. Scott 1992). If we assume that N_{H}

is a constant parameter, the differences observed between runs performed with and without tying N_{H} are difficult to reconcile. At least we can say that the sufficiency of kernel settings are an important consideration in this case, where the posterior distributions from NICER and IXPE are very disjoint. The reason is that the overlap region – where the joint posterior gets support – may be less accurately captured by the KDE approximation. At the same time, the strong disagreement between the two independent N_{H} constraints could also reflect genuine astrophysical variability in the absorbing column density. If so, N_{H} should not be assumed to be constant between the two observations, and forcing a shared value would bias the combined inference.

In combination method 2, we used the posteriors of the constant parameters of a first analysis to update the priors of the second analysis. This was done both ways (IXPE first or NICER first) to check mutual consistency. Because the NICER data are much more constraining than IXPE data (due to the larger number of counts), we see that NICER-informed IXPE posteriors strongly resemble the priors, while this is not the case when the analysis ordering is reversed. In theory we should obtain the same results, but we find that the order matters, so this method is not robust. Possibly a higher sampling resolution would converge the two results, and whether this is true could be tested with real data or synthetic data. This would likely have a high computational cost, however. We also found that the assumption of whether N_{H} is a constant parameter influences the inference of R_{eq} .

5.3 Recommendations for future work

Based on the discussion so far, we will discuss here recommendations for future work. This section is subdivided into four areas of possible improvement that we consider: the computational efficiency, model improvements (to improve physical realism), improvements to the prior, and the methodology of joint data analysis.

5.3.1 Computational speed-up

Given the computational cost of this type of analysis, it is sensible to invest in improving the computational efficiency. Another important caveat is that, as shown by Fig. 1, the pulse profiles were averaged over a period of time within which the flux was changing substantially. As we have seen in Appendix A, accounting for variable parameters increases the computational cost substantially; this was studied in more detail by Y. Kini et al. (2024a, b) within the context of TBOs.

The model evaluation of AMPs is in principle more computationally expensive compared to RMPs because of the atmosphere model, accretion disc, and polarized data. The cost here has been kept low primarily by the lack of exploration of more complex hotspot models. A straightforward way to reduce computational cost is to fine-tune the various resolution parameters and study the approximations used in X-PSI, so as to achieve the right balance of accuracy and speed, potentially tailored for the data set at hand (see also M. AlGendy & S. M. Morsink 2014; D. Choudhury et al. 2024; B. Dorsman et al. 2025; L. Mauviard et al. 2025). More drastic but potentially rewarding on the longer term would be the porting of X-PSI to graphics processing units, similar to the implementation by T. Bronzwaer et al. (2018) and T. Zhou & C. Huang (2026), the usage of machine learning to compute pulse profiles (e.g. C.-k. Chan, D. Psaltis & F. Özel 2013; P. G. Waldrop,

D. Psaltis & T. Zhao 2025) or the usage of simulation-based inference to simulate posteriors (see e.g. K. Cranmer, J. Brehmer & G. Louppe 2020, for a review).

The usage of a joint model to simultaneously fit data is an attractive strategy: it consists of only a single sampling analysis, and we attempted this in Appendix A. Because the observations were disjoint, we had to account for variable parameters, leading to high computational cost. As a consequence, we were limited to the usage of the ST-S model (i.e. antipodal hotspots) to analyse the data sets, and found significantly worse fits compared to the analyses in the main body of the paper. In addition to the measures to reduce computational cost listed so far, one could alleviate the cost here by assuming variable but apparently consistent parameters to be constant. This consistency can be based on the results of other analyses, such as how the parameters θ_p and ζ_p , were found to be in good agreement across Table 2. This will not always be an option, though, and in general we recommend restricting the analyses with joint models to (mostly) simultaneously observed data. In addition to that, the flux and pulse properties during the data selections should be preferably constant so that the reasonable assumption can be made that related parameters are constant, which may be achieved using data cuts. This reduces the computational cost by limiting the number of variable parameters, and could possibly make ST-U computationally viable.

5.3.2 Improvements to the model

There are many possible improvements to the model that could enhance its realism and therefore enhance the accuracy by which parameters are inferred. One major concern mentioned in Section 5.2 was that certain emission components could be missing, and we recommend in particular improving the modelling of the accretion environment. For example, we recommend the inclusion of the extended accretion column structure, which V. Ahlberg et al. (2024) have shown can significantly modify the hotspot emission pattern. Furthermore, although the accretion disc model has been upgraded to block radiation from the hotspots, the disc emission model is still a basic multicolour blackbody. A step up would be a general relativistic disc model (which also includes polarized emission), such as was described by V. Loktev, A. Veledina & J. Poutanen (2022). In addition, spectral modelling by e.g. C. Malacaria et al. (2025) and Z. Li et al. (2025) has shown that disc reflection modelling is an important feature of the broad-band spectrum of SRGA J144459.2–604207, so its inclusion in PPM is certainly warranted. Besides the modelling of the accretion environment, it was mentioned in Section 5.1 that the circular hotspot shapes used in this work are too simplistic compared to results from simulations. Potential improvements to the modelling of hotspot shapes for AMPs was discussed in more detail in sect. 6.4 of B. Dorsman et al. (2025).

As also stated in Section 5.1, many parameters are inferred near their prior bounds. B. Dorsman et al. (2026) analysed NICER data of SAX J1808.4–3658 with a two hotspot model and accretion disc model, and found the same trend. When they used their STU-Marg model (which features non-pulsed component that is flexible but non-physical) M , D , and i moved away from the boundary edges. We disfavour the usage of the STU-Marg however, because it introduces an arbitrary choice of bounds to the non-pulsed component that shifted the inferred parameters for B. Dorsman et al. (2026). Instead, we recommend a

physics-based model for the disc emission or at least physics-based bounds if a phenomenological model is used.

5.3.3 Improvements to the priors

Refining priors is another way of improving the accuracy and efficiency of analyses, meaning that parameter space is (accurately) narrowed down before analysis starts. As mentioned above, many parameters are inferred near their prior boundaries, which could be a hint that parameter space outside these prior distributions could lead to better fits to the data. The parameters i and d (see Section 2.3) are examples for which the prior limits were not very robust, being based only on the distribution of values found in the current literature. These could be refined via analysis of complementary data, for example through studies of X-ray bursts (see e.g. D. K. Galloway et al. 2024).

M and R_{eq} were also inferred to be close to their prior bounds in some cases. We used M and R_{eq} prior bounds that were derived from an EoS-informed prior, but not yet the actual multidimensional prior density distributions themselves. We recommend this for future work, as it would let one apply the best current prior knowledge of M and R_{eq} from theory, experiment and other astrophysical data to the inference of AMP properties (M. Hoogkamer et al. 2026). SRGA J144459.2–604207 is also a ‘clocked’ burster, for which there is a possible avenue to constrain M and R_{eq} based on the relationship between accretion rate and burst recurrence. Based on this, T. Takeda et al. (2025), suggested a mass above $2.0 M_\odot$. As a final point, we also suggest narrowing the overall parameter space down further by self-consistently linking the stellar magnetic field and accretion rate to the inner disc radius and hotspot positions, shapes, and sizes (see e.g. T. Di Salvo & A. Sanna 2022; B. Dorsman et al. 2025, and references therein).

5.3.4 Future analysis of joint data

B. Dorsman et al. (2025) and T. Salmi et al. (2025) both created synthetic data with the same parameters and used the data to test PPM with a single instrument, but they did not yet test parameter recovery with the data sets jointly. In our analysis of the NICER and IXPE data, we found that results differ significantly depending on what combination method was applied. The technique of doing joint analysis thus deserves further testing. One natural recommendation is to test the combination methods with synthetic data to understand the trade-off between accuracy and computational cost.

We also recommend the combination of low and high energies, and polarized data, because of the synergy of probing different physics in AMPs. For example, lower energy X-rays constrain the blackbody emission from the hotspots and accretion disc, whereas higher energy X-rays constrain the Comptonization of radiation at the hotspots as well as reflection off the accretion disc (including the iron line at 6.4 keV). For the outburst of SRGA J144459.2–604207 studied in this paper, it might therefore be interesting to include the ART-XC data in joint analysis. ART-XC is sensitive at higher energy (up to 35 keV) compared to NICER (2–10 keV). To add to this point, we inferred two different model fits for the plateau-like interpulse feature, but both are different compared to the proposed fit by S. V. Molkov et al. (2024). This warrants further investigation for example with ART-XC data that are partially overlapping with NICER. We note that one could

also loop in *IXPE*, *XMM–Newton*, *NuSTAR*, and *HXMT* observations, which are (partially) overlapping as well. With regards to future observations, unfortunately, *NICER* is currently non-active.⁹ However, *eXTP* (S.-N. Zhang et al. 2025) and *NewAthena* (M. Cruise et al. 2025) will be well positioned to make observations of this source. Both will have a broader energy band and much larger effective areas compared to their predecessors. *eXTP* will also have a polarimetric collector on board, with larger effective area than *IXPE*, enabling simultaneous collection of polarimetric and broad-band data. This is expected to improve constraints, with some preliminary explorations using synthetic *eXTP* data reported in A. Li et al. (2025).

6 CONCLUSION

We fitted the *NICER* and *IXPE* pulse profiles of the 2024 outburst SRGA J144459.2–604207 with an NS and accretion disc model, for the first time including polarization data. We analysed the two data sets first separately and then jointly, which was done in three different ways to verify robustness. First, we combined the posterior distributions that resulted from the separate analysis. Second, we inserted posterior distributions of constant (i.e. shared) parameters resulting in one analysis as priors into the next. That was done in both directions to test robustness. Third, in Appendix A, we used a joint model to simultaneously fit both data sets. Due to high computational cost, simplified identical antipodal hotspots were adopted there, which unfortunately proved insufficient to fit the data.

We consistently inferred a high NS mass at 1.9–2.2 $M_{\odot}$ but found two possible modes for the radius, either at 9.7–13.1 or at 13.6–14 km. For these results it is important that we note that both the mass and radius inferences may not be robust because they touch the upper bound of their priors (2.2 $M_{\odot}$ and 14 km, respectively), and we recommended further analyses with other priors and improved models. Furthermore, we inferred an NS with a cooler, large primary hotspot that is centred close to the rotational axis, which produces the bulk of the non-pulsed radiation. The secondary hotspot is hotter and is the main contributor to the pulse profile shape.

Because many of the inferred parameters were near the edges of their respective priors, we are concerned that some physics is missing in the model (such as disc reflection or reprocessing), which could cause parameter inferences to be biased. To examine such a potential bias we recommend re-analysis with improved models and with more conservative or physics-informed priors where possible.

There were notable shifts in the inferred parameters between the combination methods. Therefore the results are not fully robust, and we recommend testing with synthetic data to better understand the trade-off between accuracy and computational cost of each method. We also recommend improving the computational efficiency of the analysis and considering only simultaneously observed data to make the third combination method computationally viable. We expect that the combination of low and high energies, and polarized data will lead to complementary constraints because these probe different physics.

ACKNOWLEDGEMENTS

BD and ALW acknowledge support from European Research Council (ERC) Consolidator grant no. 865768 AEONS (PI: Watts). TS acknowledges funding by the Research Council of Finland grant no. 368807. MN is a Fonds de Recherche du Québec–Nature et Technologies (FRQNT) postdoctoral fellow. TS, AB, and JP were supported by the Research Council of Finland Centre of Excellence in Neutron-Star Physics (grants 374063 and 374064). MH, YK, and ALW acknowledge support from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) grant ENW-XL OCENW.XL21.XL21.038, Probing the phase diagram of Quantum Chromodynamics (PI: Watts). SG acknowledges the support of the Centre national d’études spatiales (CNES) and of the Agence Nationale pour la Recherche through grant ANR-25-CE31-7901-01 (ANR DENSER). CM and AP acknowledge support from the Istituto Nazionale di Astrofisica (INAF) (Research grant FANS and PULSE-X), the Italian Ministry of University and Research (PRIN MUR 2020, grant 2020BRP57Z, GEMS), and Fondazione Cariplo/Cassa Depositi e Prestiti (grant 2023–2560). This work was supported in part by the National Aeronautics and Space Administration (NASA) through the *NICER* mission. We acknowledge NWO for providing access to Snellius, hosted by SURF through the Computing Time on National Computer Facilities call for proposals. This work used the Dutch national e-infrastructure with the support of the SURF Cooperative using grant no. EINF-11691 and is subsidized by NWO Domain Science. This research was supported by the International Space Science Institute (ISSI) in Bern, through International Team project 25–657 ‘Polarimetric Insights into Extreme Magnetism’. We acknowledge extensive use of NASA’s Astrophysics Data System (ADS) Bibliographic Services and the ArXiv.

DATA AVAILABILITY

A full reproduction package is available at [10.5281/zenodo.18376256](https://doi.org/10.5281/zenodo.18376256)

REFERENCES

- Ahlberg V., Poutanen J., Salmi T., 2024, *A&A*, 682, A60
- AlGendy M., Morsink S. M., 2014, *ApJ*, 791, 78
- Baldini L. et al., 2022, *SoftwareX*, 19, 101194
- Balocco C. et al., 2025, *Astron. Telegram*, 17369, 1
- Bhattacharyya S., 2022, in Bhattacharyya S., Papitto A., Bhattacharyya D. eds, *ASSL Vol. 465, Millisecond Pulsars*. Springer International Publishing, Berlin, p. 125
- Bilous A. V. et al., 2019, *ApJ*, 887, L23
- Bobrikova A., Loktev V., Salmi T., Poutanen J., 2023, *A&A*, 678, A99
- Bogdanov S. et al., 2019, *ApJ*, 887, L26
- Bogdanov S. et al., 2021, *ApJ*, 914, L15
- Bronzwaer T., Davelaar J., Younsi Z., Mościbrodzka M., Falcke H., Kramer M., Rezzolla L., 2018, *A&A*, 613, A2
- Buchner J. et al., 2014, *A&A*, 564, A125
- Bult P. et al., 2019, *ApJ*, 885, L1
- Chan C.-k., Psaltis D., Özel F., 2013, *ApJ*, 777, 13
- Chatziioannou K., Cromartie H. T., Gandolfi S., Tews I., Radice D., Steiner A. W., Watts A. L., 2025, *Rev. Mod. Phys.*, 97, 045007
- Choudhury D. et al., 2024, *ApJ*, 971, L20
- Cranmer K., Brehmer J., Louppe G., 2020, *Proc. Natl. Acad. Sci.*, 117, 30055
- Cromartie H. T. et al., 2020, *Nat. Astron.*, 4, 72
- Cruise M. et al., 2025, *Nat. Astron.*, 9, 36

⁹See update from 2025 August 5 at <https://www.nasa.gov/missions/station/nicer-status-updates/>

- D'Antona F., Tailo M., 2022, in Bhattacharyya S., Papitto A., Bhattacharya D. eds, *ASSL Vol. 465, Millisecond Pulsars*. Springer International Publishing, Berlin, p. 201
- Das P., Salmi T., Davelaar J., Porth O., Watts A., 2025, *ApJ*, 987, 34
- Di Marco A. et al., 2022, *AJ*, 163, 170
- Di Marco A. et al., 2023, *AJ*, 165, 143
- Di Salvo T., Sanna A., 2022, in Bhattacharyya S., Papitto A., Bhattacharya D. eds, *ASSL Vol. 465, Millisecond Pulsars*. Springer International Publishing, Berlin, p. 87
- Dittmann A. J. et al., 2024, *ApJ*, 974, 295
- Dorsman B. et al., 2025, *MNRAS*, 538, 2853
- Dorsman B., Salmi T., Watts A. L., Ng M., Bobrikova A., Loktev V., Poutanen J., Wilms J., 2026, *MNRAS*, 545, staf1983
- Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601
- Feroz F., Hobson M. P., Cameron E., Pettitt A. N., 2019, *Open J. Astrophys.*, 2, 10
- Fonseca E. et al., 2021, *ApJ*, 915, L12
- Fu T. et al., 2025, *ApJ*, 980, 161
- Galloway D. K. et al., 2020, *ApJS*, 249, 32
- Galloway D. K., Goodwin A. J., Hilder T., Waterson L., Cupák M., 2024, *MNRAS*, 535, 647
- Gandolfi S., Carlson J., Reddy S., 2012, *Phys. Rev. C*, 85, 032801
- Gendreau K. C. et al., 2016, in den Herder J.-W. A., Takahashi T., Bautz M. eds, *Proc. SPIE Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*. SPIE, Bellingham, p. 99051H,
- Gierliński M., Poutanen J., 2005, *MNRAS*, 359, 1261
- Harding A. K., Muslimov A. G., 2001, *ApJ*, 556, 987
- Hebeler K., Holt J. D., Menéndez J., Schwenk A., 2015, *Annu. Rev. Nucl. Part. Sci.*, 65, 457
- Heinke C. O., Zheng J., Maccarone T. J., Degenaar N., Bahramian A., Sivakoff G. R., Toor S., 2025, *ApJS*, 279, 57
- Hoogkamer M., Kini Y., Salmi T., Watts A. L., Buchner J., 2025, *Phys. Rev. D*, 112, 023008
- Hoogkamer M. et al., 2026, *Phys. Rev. D*, 113, 063049
- Ibragimov A., Poutanen J., 2009, *MNRAS*, 400, 492
- Jakab Z., Morsink S. M., 2025, *ApJ*, 994, 163
- Kajava J. J. E., Ibragimov A., Annala M., Patruno A., Poutanen J., 2011, *MNRAS*, 417, 1454
- Kass R. E., Raftery A. E., 1995, *J. Am. Stat. Assoc.*, 90, 773
- Kini Y. et al., 2023, *MNRAS*, 522, 3389
- Kini Y. et al., 2024a, *MNRAS*, 527, 8118
- Kini Y. et al., 2024b, *MNRAS*, 535, 1507
- Kini Y. et al., 2025, *MNRAS*, 541, 46
- Kini Y. et al., 2026, *ApJ*, 1005, 201
- Kislat F., Clark B., Beilicke M., Krawczynski H., 2015, *Astropart. Phys.*, 68, 45
- Koehn H. et al., 2025, *Phys. Rev. X*, 15, 021014
- Kulkarni A. K., Romanova M. M., 2013, *MNRAS*, 433, 3048
- LaMarr B., Prigozhin G., Remillard R., Malonis A., Gendreau K. C., Arzoumanian Z., Markwardt C. B., Baumgartner W. H., 2016, in den Herder J.-W. A., Takahashi T., Bautz M. eds, *Proc. SPIE Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*. SPIE, Bellingham, p. 99054W,
- Lasota J.-P., 2001, *New Astron. Rev.*, 45, 449
- Lattimer J. M., 2021, *Annu. Rev. Nucl. Part. Sci.*, 71, 433
- Leahy D. A., Morsink S. M., Cadeau C., 2008, *ApJ*, 672, 1119
- Leahy D. A., Morsink S. M., Chung Y.-Y., Chou Y., 2009, *ApJ*, 691, 1235
- Li A. et al., 2025, *Sci. China Phys. Mech. Astron.*, 68, 119503
- Li Z. et al., 2025, *ApJ*, 990, 15
- Loktev V., Veledina A., Poutanen J., 2022, *A&A*, 660, A25
- Luo J. et al., 2021, *ApJ*, 911, 45
- Makishima K., Maejima Y., Mitsuda K., Bradt H. V., Remillard R. A., Tuohy I. R., Hoshi R., Nakagawa M., 1986, *ApJ*, 308, 635
- Malacaria C. et al., 2025, *A&A*, 699, A288
- Mandal M., Naik S., Chhotaray B., 2025, *MNRAS*, 542, 443
- Mao Y.-H., Dorsman B., Watts A. L., Salmi T., Poutanen J., Li X.-D., 2026, preprint ([arXiv:2606.14493](https://arxiv.org/abs/2606.14493))
- Matsuoka M. et al., 2009, *PASJ*, 61, 999
- Mauviard L. et al., 2025, *ApJ*, 995, 60
- Miller M. C., Lamb F. K., 1998, *ApJ*, 499, L37
- Miller M. C. et al., 2019, *ApJ*, 887, L24
- Miller M. C. et al., 2021, *ApJ*, 918, L28
- Miller M. C. et al., 2026, *ApJ*, 1000, L48
- Mitsuda K. et al., 1984, *PASJ*, 36, 741
- Molkov S. V. et al., 2024, *A&A*, 690, A353
- Morsink S. M., Leahy D. A., 2011, *ApJ*, 726, 56
- Morsink S. M., Leahy D. A., Cadeau C., Braga J., 2007, *ApJ*, 663, 1244
- Ng M. et al., 2024, *ApJ*, 968, L7
- Oliva G. A., Frutos-Alfaro F., 2021, *MNRAS*, 505, 2870
- Papitto A. et al., 2025, *A&A*, 694, A37
- Patruno A., Watts A. L., 2021, in Belloni T. M., Méndez M., Zhang C. eds, *ASSL Vol. 461, Timing Neutron Stars: Pulsations, Oscillations and Explosions*. Springer, Berlin, p. 143, preprint ([arXiv:1206.2727](https://arxiv.org/abs/1206.2727))
- Pechenick K. R., Ftaclas C., Cohen J. M., 1983, *ApJ*, 274, 846
- Poutanen J., 2020, *A&A*, 641, A166
- Poutanen J., Beloborodov A. M., 2006, *MNRAS*, 373, 836
- Poutanen J., Gierliński M., 2003, *MNRAS*, 343, 1301
- Poutanen J., Ibragimov A., Annala M., 2009, *ApJ*, 706, L129
- Prigozhin G. et al., 2016, in den Herder J.-W. A., Takahashi T., Bautz M. eds, *Proc. SPIE Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*. SPIE, Bellingham, p. 99051I,
- Qi L. et al., 2025, *ApJ*, 981, 99
- Reardon D. J. et al., 2024, *ApJ*, 971, L18
- Remillard R. A. et al., 2022, *AJ*, 163, 130
- Riffert H., Meszaros P., 1988, *ApJ*, 325, 207
- Riley T. E., Raaijmakers G., Watts A. L., 2018, *MNRAS*, 478, 1093
- Riley T. E. et al., 2019, *ApJ*, 887, L21
- Riley T. E. et al., 2021, *ApJ*, 918, L27
- Riley T. E. et al., 2023, *J. Open Source Softw.*, 8, 4977
- Romanova M. M., Ustyugova G. V., Koldoba A. V., Lovelace R. V. E., 2004, *ApJ*, 610, 920
- Rutherford N. et al., 2024, *ApJ*, 971, L19
- Salmi T., Nättälä J., Poutanen J., 2018, *A&A*, 618, A161
- Salmi T., Loktev V., Korsman K., Baldini L., Tsygankov S. S., Poutanen J., 2021, *A&A*, 646, A23
- Salmi T. et al., 2022, *ApJ*, 941, 150
- Salmi T. et al., 2023, *ApJ*, 956, 138
- Salmi T. et al., 2024a, *ApJ*, 974, 294
- Salmi T. et al., 2024b, *ApJ*, 976, 58
- Salmi T. et al., 2025, *MNRAS*, 538, 2562
- Scott D., 1992, *Multivariate Density Estimation: Theory, Practice, and Visualization*. A Wiley-Interscience Publication, Wiley, Hoboken, https://books.google.nl/books?id=7crCUS_F2ocC
- Silva H. O., Pappas G., Yunes N., Yagi K., 2021, *Phys. Rev. D*, 103, 063038
- Skilling J., 2004, in Fischer R., Preuss R., Toussaint U. V. eds, *AIP Conf. Ser. Vol. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*. AIP, Melville, NY, p. 395
- Soffitta P. et al., 2021, *AJ*, 162, 208
- Sunyaev R. et al., 2021, *A&A*, 656, A132
- Takeda T. et al., 2025, *PASJ*, 77, L24
- Tolos L., Fabbietti L., 2020, *Prog. Part. Nucl. Phys.*, 112, 103770
- Viironen K., Poutanen J., 2004, *A&A*, 426, 985
- Vinciguerra S., Salmi T., Watts A. L., Choudhury D., Kini Y., Riley T. E., 2023, *ApJ*, 959, 55
- Vinciguerra S. et al., 2024, *ApJ*, 961, 62
- Waldrop P. G., Psaltis D., Zhao T., 2025, preprint ([arXiv:2506.11194](https://arxiv.org/abs/2506.11194))
- Watts A. L., 2012, *ARA&A*, 50, 609
- Watts A. L. et al., 2016, *Rev. Mod. Phys.*, 88, 021001
- Weisskopf M. C. et al., 2022, *J. Astron. Telesc. Instrum. Syst.*, 8, 026002
- Zhang S.-N. et al., 2025, *Sci. China Phys. Mech. Astron.*, 68, 119502
- Zhou T., Huang C., 2026, *A&A*, 709, A111

Table A1. An overview of the analysis with the joint model and the inferred parameters. This table follows the format of Table 2, see also the caption of that table.

Name	Combination method 3
Data	NICER, <i>IXPE</i> <i>QU</i>
Hotspot config.	ST-S
Sampling aspects	
Num. params.	23
Samp. eff. (X-PSI)	0.1
Run cost (kch)	115.2
Live points	1000
$\ln(\mathcal{Z})$	$-1.5955e + 5$
NICER fit quality	
χ^2/DOF	30744/28777
$\hat{\chi}^2$	1.0683
p -value	5.268e-16
<i>IXPE</i> fit quality	
χ^2/DOF	5251/4825
$\hat{\chi}^2$	1.0888
p -value	1.1861e-5
M [$M_\odot$]	$1.015^{+0.024}_{-0.011}$
R_{eq} [km]	$13.90^{+0.08}_{-0.16}$
D [kpc]	$3.13^{+0.09}_{-0.09}$
i [deg]	$50.25^{+0.41}_{-0.19}$
N_{H} [10^{21}cm^{-2}]	$25.52^{+0.06}_{-0.06}$
$\phi_{\text{p},\text{N}}$ [cycles]	$0.056^{+0.003}_{-0.003}$
$\theta_{\text{p},\text{N}}$ [deg]	$6.9^{+1.0}_{-0.8}$
$\zeta_{\text{p},\text{N}}$ [deg]	$86.42^{+0.17}_{-0.21}$
$T_{\text{seed},\text{p},\text{N}}$ [keV]	$0.812^{+0.012}_{-0.013}$
$T_{\text{e},\text{p},\text{N}}$ [keV]	$21.3^{+0.7}_{-0.6}$
$\tau_{\text{p},\text{N}}$ [—]	$2.40^{+0.03}_{-0.04}$
$\phi_{\text{p},\text{I}}$ [cycles]	$-0.319^{+0.004}_{-0.004}$
$\theta_{\text{p},\text{I}}$ [deg]	$24.0^{+2.5}_{-2.5}$
$\zeta_{\text{p},\text{I}}$ [deg]	$84.2^{+0.4}_{-0.5}$
$T_{\text{seed},\text{p},\text{I}}$ [keV]	$0.568^{+0.007}_{-0.007}$
$T_{\text{e},\text{p},\text{I}}$ [keV]	$46.2^{+5.3}_{-4.5}$
$\tau_{\text{p},\text{I}}$ [—]	$1.30^{+0.06}_{-0.05}$
$T_{\text{in},\text{N}}$ [keV]	$0.54^{+0.01}_{-0.01}$
$R_{\text{in},\text{N}}$ [km]	$14.46^{+0.13}_{-0.17}$
α_{NICER} [—]	$0.811^{+0.019}_{-0.009}$
α_{DU2} [—]	$1.0218^{+0.0018}_{-0.0018}$
α_{DU3} [—]	$0.9939^{+0.0017}_{-0.0017}$
χ_0 [rad]	$-0.7^{+1.8}_{-0.6}$

APPENDIX A: COMBINATION METHOD 3 (SIMULTANEOUS FIT)

This appendix presents the results of the analysis using a joint model to simultaneously fit the NICER and *IXPE* data. Because the observations of the data sets were disjoint, two instances of all variable parameters are used. For some parameters, that may turn out to be unnecessary if we end up inferring similar values. However, we choose to always treat the two instances as independent parameters because we do not know a priori how




Figure A1. Same figure as Fig. 4, but for the joint model analysis. In this case, the left panels show the projection, spectra, pulse profiles, and residuals associated with the model during the NICER observation, and the right panels with the *IXPE* (DU₁) observation. DU₂ and DU₃ are not shown, but they are very similar to DU₁. See the caption of Fig. 4 for more detail on all the components in the plot.

different the inferred values will be. For this reason – and because the emission is computed twice (for NICER and for *IXPE*) – the analysis has significantly higher computational cost. To get a run to complete within the budget, we reduced the model complexity from ST-U to ST-S (defined in Section 2.1). This brings the total parameters down from a potential 35 to 23. Still, at 115 kch, this run was significantly more expensive than even the *IXPE* and NICER runs added together, as shown in Table 2 and Table A1.

When we look at the inferred parameters, shown in Figure A2, we note that the inferred M is now towards the lower bound (rather than upper bound) of the prior, at $M = 1.015^{+0.024}_{-0.011} M_\odot$. On the other hand, the inferred R_{eq} is consistent with high radius results preferred by the ST-U results that analyse NICER data first, at $13.90^{+0.08}_{-0.16}$ km. D and i tend towards the lower bound of their priors, which is a similar finding to all previous results. N_{H} , at $(25.52 \pm 0.06) \times 10^{21} \text{ cm}^{-2}$, is roughly in the middle of the inferred values from the NICER and *IXPE* analyses. We also find again that the decrease in flux is explained mainly by a decrease of $T_{\text{seed},\text{p}}$, consistent with the idea that the accretion rate decreases.

It is also interesting to note that the inner disc radius inferred for the NICER data is very small in this analysis, much smaller compared to the ST-U findings, and blocks almost the entire




Figure A2. Posterior distributions of the inferred constant parameters of the joint model analysis (combination method 3, see Section A). See the caption of Fig. 5 for a detailed description of the plot elements.

Southern hemisphere and secondary hotspot. This finding is not consistent with the classic AMP picture where accreted material from the inner accretion disc that truncates some way from the stellar surface is funnelled towards the magnetic poles of the NS, forming hotspots, but instead looks more like boundary layer accretion. The hotspots are inferred to be large during both the NICER and *IXPE* observations, but the spot colatitudes have increased significantly between the two observations.

When we look at the fit statistics, we see that the joint model fits of either data set are significantly worse than the fits in the main body of the paper. The χ^2/DOF for the NICER data is $30772/28777 = 1.0693$, corresponding to a p -value of 2.0629×10^{-16} , and the χ^2/DOF for the Stokes I data of the three *IXPE* detectors is $5250/4825 = 1.0881$, corresponding to a p -value of 1.2619×10^{-5} . We see again the NICER aspects in the NICER residual. In addition, here the worse fit is manifested in a phase-dependent structure visible in both residuals, shown by Fig. A1. This indicates that the ST-S joint model is strongly disfavoured compared to the separate ST-U analyses. One caveat is that there were only 1000 live points used here, and that can be considered insufficient for an analysis with 23 parameters, especially if the likelihood surface is complex. Due to the expected high computational cost, these results were not tested for convergence against changes to the sampler settings. Another caveat in this analysis is that four parameters: M , D , $T_{e,N}$, and α_N all converged towards prior boundaries, hinting that parameter space outside current prior boundaries could lead to better fits.

To conclude, the antipodal and identical hotspots that ST-S enforces, appear a serious restriction that limits the ability of the model to fit the data. For future work, we recommend the application of the ST-U configuration instead. We describe methods to overcome the computational cost in Section 5.3.1.

APPENDIX B: EXTRA CORNER PLOTS

Figs B1–B4 show the full posterior distributions of all the ST-U analyses included in this work. This includes the *IXPE* Stokes Q and U analysis, in which the Stokes I data were not used. Figs B5 and B6 show the posterior distributions of the variable parameters (as observed by NICER or *IXPE*) resulting from combination method 3.




Figure B1. Posterior distributions of the geometry parameters for the IXPE analyses. See the parameter and model definitions in Table 1, and Sections 4.2 and 4.4.2. The dash-dotted curves represent the prior distributions (which are identical for the IQU and QU cases). The 1D shaded intervals contain 68.3 per cent of the posterior mass, and the 2D contours contain 68.3 per cent, 95.4 per cent, and 99.7 per cent of the posterior mass. The IQU posteriors of Θ_p are difficult to see, because they are at the lower bound of the prior.




Figure B2. Posterior distributions of the non-geometry parameters for the *IXPE* analyses. See the parameter and model definitions in Table 1, and Sections 4.2 and 4.4.2. The dash-dotted curves represent the prior distributions (which are identical for the *IQU* and *QU* cases). The 1D shaded intervals contain 68.3 per cent of the posterior mass, and the 2D contours contain 68.3 per cent, 95.4 per cent, and 99.7 per cent of the posterior mass.




Figure B3. Posterior distributions of the non-geometry parameters for the NICER analyses. See the parameter and model definitions in Table 1 and Section 2.1. The dash-dotted curves represent the prior distributions (which only differ in N_{H}). The 1D shaded intervals contain 68.3 per cent of the posterior mass, and the 2D contours contain 68.3 per cent, 95.4 per cent, and 99.7 per cent of the posterior mass.




Figure B4. Posterior distributions of the geometry parameters for the NICER analyses. See the parameter and model definitions in Table 1, and Section 2.1. The dash-dotted curves represent the prior distributions (which only differ in N_{H}). The 1D shaded intervals contain 68.3 per cent of the posterior mass, and the 2D contours contain 68.3 per cent, 95.4 per cent, and 99.7 per cent of the posterior mass.




Figure B5. Posterior distributions of the variable parameters during the NICER observation for the joint model analysis. See the parameter and model definitions in Table 1, and Section 2.1. The dash-dotted curves represent the prior distributions (which only differ in N_H). The 1D shaded intervals contain 68.3 per cent of the posterior mass, and the 2D contours contain 68.3 per cent, 95.4 per cent, and 99.7 per cent of the posterior mass.




Figure B6. Posterior distributions of the variable parameters during the *IXPE* observation for the joint model analysis. See the parameter and model definitions in Table 1, and Section 2.1. The dash-dotted curves represent the prior distributions (which only differ in N_H). The 1D shaded intervals contain 68.3 per cent of the posterior mass, and the 2D contours contain 68.3 per cent, 95.4 per cent, and 99.7 per cent of the posterior mass.

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.