WIDEBAND OBSERVATIONS OF RADIO PULSARS

Timothy Thomas Pennucci
White Plains, New York

B.A. Astrophysics, Columbia University in the City of New York, 2009

M.S. Astronomy, University of Virginia, 2011

A Dissertation Presented to the Graduate
Faculty of the University of Virginia
in Candidacy for the Degree of
Doctor of Philosophy

Department of Astronomy

University of Virginia
August 2015

Committee:

Professor Scott M. Ransom
Dr. Paul B. Demorest
Professor D. Mark Whittle
Professor Phil L. Arras
Professor R. Craig Group

Abstract

Pulsars are exotic objects which have yielded a bounty of important astrophysical results. As rapidly rotating, highly magnetized neutron stars, pulsars' stable rotation and beamed radio emission enables their use as interstellar laboratory clocks. The extraordinary timing regularity of the millisecond pulsar (MSP) population permits some of the most precise measurements in astronomy. The discovery of MSPs raised the probability of directly detecting gravitational waves for the first time. Ongoing efforts by several pulsar timing array (PTA) collaborations compliment the ground- and space-based efforts of laser interferometers. One such PTA is the North American Nanohertz Observatory for Gravitational Waves (NANOGrav). NANOGrav has recently employed a new set of wideband instruments to increase the sensitivity of their PTA, and the future of pulsar astronomy is moving towards progressively larger bandwidths. In this dissertation, we address the benefits and issues from adopting the new instrumentation, particularly for the scientific motivations of NANOGrav. We first develop a measurement technique for simultaneously obtaining pulse times-of-arrival (TOAs) and dispersion measures (DMs) using 2D models of evolving Gaussian components. We then apply the methodology broadly to a variety of pulsars, including a bright, test MSP in a globular cluster, the Galactic Center magnetar, and the entire suite of 37 MSPs from the NANOGrav 9-year data set. For a subset of these MSPs, we make targeted observations at specific orbital phases aimed at improving the timing models and constraining the Shapiro delay. With a few exceptions, we find positive or consistent timing results from the implementation of our first generation wideband timing protocol. Some highlights include: improved measurement uncertainties, mitigation of chromatic ISM effects, a reduction in the number of timing parameters and TOAs, signs of chromatic DMs, and at least one new pulsar mass.

To Annine and James,

because my story needs

a heroine and a foil.

Acknowledgments

I would be absolutely remiss if I did not point out the Authentic Science Research Program at my high school, WPHS. Without this program, I may have taken a much more random walk to scientific research. ASR led me to Charles Liu, the AMNH, the NRAO, Columbia, UVa...but it first led me to my guidance counselor's office, where I filled a line on a questionnaire that asked what level of education I wished to obtain with the words "PhD in Astrophysics". I am fortunate to have been part of such an (unfortunately) exceptional public school system, and I wholeheartedly thank all my teachers for their often unappreciated service.

Charles Liu deserves a further acknowledgement. Had Charles not been so kind as to guide me along for two years and let me sit among professional astronomers, graduate students, and REU students as a teenager in the Perkin Reading Room of the American Museum of Natural History, I would not be here today. Of this much I am sure.

David Helfand, Frits Paerels, and Fernando Camilo at Columbia were integral in launching my graduate career. David and Frits, in particular, inspired my return to Morningside.

I have been an active participant in NANOGrav for my entire graduate career and I owe all of its members many thanks. I have benefited tremendously from the expertise of this constellation of scientists, particularly those of the timing group. Furthermore, NANOGrav enabled me to travel and interact with astronomers from other PTA collaborations around the world, and just for this I am immensely obliged. I must also single out Justin Ellis for his role in helping me tie up some of the loose ends of my research as time ran out.

Much appreciation also is shown for my immediate academic community, the UVa

Department of Astronomy and the National Radio Astronomy Observatory — faculty, astronomers, staff, and grads alike — for making this possible. I have gained much mileage from the NRAO, as far back as high school. Of course, I have to highlight the Ransom Factory: Scott Ransom, Paul Demorest, Ryan Lynch, Jintao Luo, Anya Bilous, Siraprapa Sanpa-Arsa (Tuck!), Brian Prager, and Thankful Cromartie.

A resounding acknowledgment goes out to Anya Bilous, my friend and collaborator, with whom I have an unresolved bet regarding the future detection rate of FRBs (I cannot recall the exact details, but when CHIME makes some detections, she owes me a bottle of bourbon). Thanks for all the tea, the language, and the ride to Dulles.

I am indebted to my advisors Scott Ransom and Paul Demorest beyond recognition. In addition to training me with a wealth of knowledge and skills pertaining to pulsar astronomy, Scott has been good for everything from showing me how to untar a gzipped archive, to beating me in pull-up contests (on a moving bus, no less), to enabling my life as a globetrotting romantic. As for Paul, since we both first showed up in Charlottesville in 2007, he as a Jansky Fellow and I as an REU student, he has been indispensable. There is a reason why some of us in the Ransom Factory count work-related tasks in units of Demorests. In just about every realm of my research where I have run into trouble, I could count on Paul to not only be of assistance but to make it look easy, in which case we say one was #demorested. Infinite thanks to you both.

There are no satisfactory words with which I can thank my family, so I will say only that I hope I can make them proud. Because this marks both the end and beginning of certain things, then to my dear woman, Emőke Csernus, I say köszönöm szépen, hogy benne vagyunk. Meet me at the top!

Preface

One day, Charles Liu spelled out a URL for me: `www.nrao.edu`. I have been of the radio persuasion ever since. I took only a few steps in my walk to radio pulsars in particular, and this document embodies my initial contribution to the field. If a plot is worth a thousand words, then consider this document unabashedly verbose; I am unapologetic about including useful figures for all of the pulsars incoporated in this study.

I have been fortunate since the beginning of my studies of pulsars to have sampled most areas of the P-$\dot{P}$ diagram. This dissertation is primarily concerned with millisecond pulsars, as well as one magnetar, but I also get to claim discovery for a single slow pulsar in the GBT 350 MHz Drift-scan Survey (Boyles et al. 2013; Lynch et al. 2013). In the beginning, as an undergraduate REU student at the NRAO, I searched through some ten thousand(s) candidate pulsars and *finally* identified but one slow, canonical pulsar, the 426-ms pulsar J1612+2008[1]. Let this be a lesson to new students who are enticed by wily, animated pulsar astronomers quoting discovery rates!

Along with the document before you, this dissertation has associated with it more than 22k/5k words/lines of finished `python` code used for the bulk of the work, most of which has been rewritten several times. All of the work in this document was supported by the NANOGrav PIRE grant (NSF 0968296) and has resulted in at least two publications.

Enjoy.

18 June 2015, Room 109

[1]a.k.a. "PSR Tacocat"; attribution by A. Jáuregui, private communication.

Table of contents

List of Figures

List of Tables

Chapter 1

Introduction to Pulsars & Pulsar Timing

"Timing is everything." —The narcissistic horologist

1.1 Opening Remarks

After nearly fifty years of study, pulsars continue to be extremely fruitful objects for scientific enquiry. A Nobel Prize was awarded for their discovery[1] at least in part because of the impact that pulsars were anticipated to have on the community. Indeed, in this half-century, the study of pulsars has contributed to astrophysical fields relating to the interstellar medium, binary evolution, plasma physics, supernovae, nuclear physics, extrasolar planets, hierarchical galaxy formation, magnetic fields, and gravity – the lattermost for which a second Nobel Prize was awarded[2]. All such contributions stem from the facts that pulsars are exotic objects and can function like laboratory clocks situated in interstellar space.

The techniques associated with making use of pulsars' clock-like behavior are collectively called "pulsar timing". Pulsar timing — particularly at radio frequencies — yields some of the most precise measurements in all of astronomy, from nearly percent-level interstellar distances to neutron star mass determinations that are precise to several times the mass of Jupiter. The history of pulsar astronomy has been paralleled by a monotonic increase in the precision obtained from pulsar timing due to the construction of bigger dishes, the design of more sensitive radio receivers, the observation of pulsars in broader bandwidths, and the improvements in signal detection and pulsar timing algorithms. One significant discontinuity in this history, however, was the discovery of the millisecond pulsar population. Because they have the shortest spin periods and the least amount of unpredictable behavior, millisecond pulsars provide the most precise timing measurements.

Most of the present work is concerned with the measurements that can be obtained

[1] A. Hewish, 1974, though the prize was not shared with his postdoctoral student and acknowledged discoverer of pulsars, Jocelyn Bell.

[2] R. Hulse & J .Taylor, 1993.

from the current generation of pulsar instruments, which are significant in this history for being the first set of instruments to process almost a gigahertz of bandwidth in real time explicitly for the study of millisecond pulsars at frequencies relevant to high-precision timing experiments. However, the advantages that come with a broader bandwidth are only achieved if there is a parallel development in pulsar timing methodologies, as will be explained and explored.

All of the advances mentioned above are leading to what could be the climax of pulsar astronomy: the direct detection of gravitational waves by means of a pulsar timing array[3]. In this dissertation, we make progress on the front of developing methods for the current and next generation of pulsar timing array experiments and their instruments, and showcase our developments in a series of cutting edge measurements of radio pulsars. Here, we give a brief overview to pulsars, millisecond pulsars, pulsar timing arrays and one of the central problems that we address in the chapters that follow.

1.2 Pulsars

The observations associated with pulsars have grown to include quite a range of phenomena since the initial detection of stable, periodic radio emission with a period of ~ 1 s (Hewish et al. 1968). A modern, working definition that would encompass all objects with the designation of "pulsar" might be: a neutron star that reveals its spin period by the combination of anisotropic emission of electromagnetic radiation and rotation. Therefore, due to the "lighthouse effect", any such rotating neutron star whose anisotropic emission crosses our line of sight might be seen to pulse and be called a pulsar. In addition to canonical and radio millisecond pulsars, this

[3] One notable alternative is the discovery of a millisecond pulsar in orbit around a black hole.

would include accretion-powered millisecond X-ray pulsars, gamma-ray-only pulsars, magnetars, rotating radio transients, and the so-called isolated neutron stars, but excludes other transient, pulsed emission (most notably, the as-yet-unidentified "fast radio bursts").

The underlying commonality, of course, is that pulsars are neutron stars, which were predicted to exist some eighty years ago (Baade & Zwicky 1934). While the original pulsars could have been explained as being white dwarf stars or radial oscillations, the discovery of the Crab pulsar (Staelin & Reifenstein 1968), with its short 33 ms spin period and intrinsic spin down, could only be explained by a rotating neutron star (Pacini 1967, 1968; Gold 1968; Comella et al. 1969; Richards & Comella 1969; Lyne & Graham-Smith 2012). Because a self-gravitating object can only spin so quickly without breaking up, a spin period of 33 ms implied a minimum density approaching nuclear matter; for an object larger than the Chandrasekhar mass ($\sim$1.4 $M_\odot$), it could be no larger than $\sim$20 km in radius. This picture is consistent with the remnants predicted to be the end products of supernovae by Baade & Zwicky (1934). Incidentally, the discovery of the Crab and Vela pulsars (Large et al. 1968) also codified the hypothesis that these neutron stars are born as a result of supernova explosions when high mass stars end their fusion powered life. Neutron stars are held up against further gravitational collapse in part due to degeneracy pressure of neutrons, but the equation of state for neutron star material remains a holy grail in nuclear physics, as it perhaps describes the ultimate state of observable dense matter in the Universe. Pulsar timing is one of the few ways in which we can make reliable and precise mass measurements of neutron stars in order to constrain the equation of state; this is a topic we return to in Chapter 5.

The angular momentum and magnetic field conferred to a neutron star as a result of its progenitor's death leaves it with a spin period on the order of tens of milliseconds and a field strength around 10^{12} G. The range of observed pulsar parameters spans almost four dex in spin period P ($\sim$0.001–10 s) and 10 dex in spin period derivative $\dot{P}$ ($\sim 10^{-20}$–10^{-10} s s^{-1}). However, most pulsars have a spin-down luminosity ($\dot{E} \approx P^{-3}\dot{P}/100\ L_{\odot}$) close to one solar luminosity, plus or minus three dex. The combination of the pulsar's spin, spin down, magnetic field, and gravitational potential drive all of the exotic observational phenomena seen, which includes: "classical" coherent, polarized radio emission, giant nanosecond-duration pulses with the largest brightness temperatures of anything observed, magnetar X-ray outbursts, and pulse mode switching on the order of a single rotation.

Despite five decades of study, most of these phenomena are still poorly understood. One standard picture describing the pulsed radio emission, since it is relevant to the observations we make in the following chapters, is as follows. It was quickly realized that a rotating neutron star with a dipolar magnetic field will not remain in a vacuum; the induced electric fields far exceed the strength of gravity and set up a screening plasma to form the magnetosphere (Goldreich & Julian 1969). Gaps in the magnetosphere provide accelerating potentials for a primary plasma. This plasma produces gamma ray photons via curvature radiation or inverse Compton scattering. In the presence of the magnetic field, an electron-proton pair production cascade ensues to amplify the original plasma density by several orders of magnitude (Sturrock 1971; Ruderman & Sutherland 1975). The enhanced density plasma is responsible for the radio emission at some distance away from the neutron star surface — either above the polar regions, or off near the light-cylinder boundary, where corotation with the neutron star surface equals the speed of light. An undetermined coherent process is then responsible for producing the observed average radio light curves, which (given the observing geometry) are often characterized by having multiple components of a wide variety of shapes, polarizations, relative positions, and spectra (cf. Figures 4.156–4.228).

A detailed overview of pulsar astronomy can be found in the standard texts Lorimer & Kramer (2005) and Lyne & Graham-Smith (2012). However, next we further introduce millisecond pulsars since they play a central role in this work.

1.2.1 Millisecond Pulsars

Of the roughly 2400 known pulsars[4], about 15% have spin periods less than $\sim$20 ms – these are the millisecond pulsars (MSPs), which are further characterized by having the smallest spin down rates ($\sim 10^{-4} - 10^{-2}$ ns yr^{-1}) and smallest magnetic field strengths ($\lesssim 10^{10}$ G). Almost all MSPs are in binary systems, where the angular momentum obtained by mass transfer onto the neutron star from the companion during the latter's evolution is responsible for the distinctly short spin period. This "recycling model" has only very recently been definitively confirmed (broadly speaking) with the observations of several systems swinging between an apparent accreting, low-mass X-ray binary state, and a "normal" radio millisecond pulsar mode (Archibald et al. 2009; Papitto et al. 2013; Roy et al. 2015)[5]. While most MSPs are found in the field[6], a large fraction ($\sim$35%) are found in globular clusters due to the old stellar populations and the increased probability of chance encounters to be spun-up[7].

The rough precision with which a pulsar can be timed is proportional to its spin frequency, making MSPs the best candidates for timing experiments. In addition to

[4]`http://www.atnf.csiro.au/people/pulsar/psrcat/`, (Manchester et al. 2005).

[5]In 2011, the author's own efforts to find this link by searching for radio pulsations from the peculiar, 600 Hz accretion-powered millisecond X-ray pulsar J00291+5934 were unsuccessful.

[6]`http://astro.phys.wvu.edu/GalacticMSPs/GalacticMSPs.txt`

[7]`http://www.naic.edu/~pfreire/GCpsr.html`

this, however, is the important fact that MSPs are much more stable; slower pulsars with larger magnetic fields are known to exhibit glitches – sudden changes in the moment of inertia that translate into an immediate increase (or decrease) in the spin frequency, followed by decay – and "timing noise". Timing noise can be modeled as the result of a random walk in either a pulsar's rotational phase, spin period, or period derivative (Shannon & Cordes 2010), and manifests as a stochastic, long-term drifting trend in the clock's behavior that ultimately limits the timing precision. The extent to which timing noise is present in MSPs is an important question for high-precision experiments, such as pulsar timing arrays, which we will introduce following the next section.

1.3 Pulsar Timing

Here, we introduce the fundamental concepts of pulsar timing, borrowing from Lorimer & Kramer (2005) and Demorest (2007), where more thorough treatments can be found. At the heart of pulsar timing is the fact that one can count rotations of the neutron star, and so each observed pulse can be assigned an integer[8]. The pulse phase ϕ can be written as a function of time t by integrating a Taylor expansion of the spin frequency f,

$$\phi(t) = \phi(t_0) + f_0(t - t_0) + \frac{1}{2}\dot{f}(t - t_0)^2 + \frac{1}{6}\ddot{f}(t - t_0)^3 + ... \tag{1.1}$$

This is the most basic timing model for a given pulsar, as seen in the reference frame of the pulsar. Generally, $\ddot{f}$ and higher terms are only measurably different from zero

[8]Hence the maxim, "Pulsar timing unambiguously accounts for each and every rotation of the neutron star."

in the cases of timing noise or large intrinsic spin-down evolution[9]. $\phi_0 = \phi(t_0)$ can be arbitrarily assigned an integer, in turn defining t_0. Thus, when the observed pulse times-of-arrival (TOAs) are transformed into the pulsar's reference frame, then those set of times $\{t_i\}$ should correspond to integer values $\phi(t_i)$.

Most of the complexity in the timing model arises from the non-inertial frame of the Earth-based observatory, the orbit of a binary pulsar, and the intervening interstellar medium (ISM). Each of these introduces a number of delays or corrections into the transformation of the observed TOAs. The parameterizations of these delays contain useful quantities, e.g., the sky position of the pulsar, its parallax (distance) and proper motion, the Keplerian orbital parameters, and the column density of free electrons along the line of sight. This latter quantity is called the dispersion measure (DM), which quantifies how much a TOA at a given frequency lags an infinite frequency TOA due to the refractive nature of the ionized ISM. The TOAs need to be accurately extrapolated to infinite frequency arrival times because a *changing* DM will introduce variable delays that are a significant source of noise in high-precision timing experiments (Lee et al. 2014).

The measurement of TOAs is covered in Chapter 2, which expands upon the standard template-matching technique to include a simultaneous measurement of the DM. Once a set of TOAs has been obtained, one bootstraps a timing model by measuring a small number of parameters with a subset of TOAs, and adds more parameters into the timing model as they become significant. The timing model should converge to a "solution", in which case one says they have "solved the pulsar" or obtained a phase-connected timing solution; this effectively would allow the observer to go long periods of time without observing the pulsar and maintain confidence that all rota-

[9]We are also ignoring non-inertial accelerations from moving in a gravitational potential.

tions are accounted for in the model, provided it is good enough[10]. The basic problem is formulated in Lorimer & Kramer (2005, §8.2.6, Equation 8.15):

$$\chi^2_{\mathrm{TOA}} = \sum_{i}^{\#\ \mathrm{TOAs}} \left(\frac{n_i - \phi(t_i)}{\sigma_{\mathrm{TOA},i}} \right)^2, \tag{1.2}$$

where n_i is the closest integer to $\phi(t_i)$ and $\sigma_{\mathrm{TOA},i}$ is the uncertainty on the ith TOA. The process of minimizing χ^2_{TOA} is synonymous with fitting a timing model with pulsar timing software such as `tempo`[11]. However, because TOAs can be correlated, it is simpler to formulate the problem in terms of a general covariance matrix, $\mathbf{\Sigma}$, which may have off-diagonal elements in addition to the heteroscedastic $\sigma^2_{\mathrm{TOA},i}$ on the diagonal:

$$\chi^2_{\mathrm{TOA}} = \mathbf{r}^{\mathrm{T}} \mathbf{\Sigma}^{-1} \mathbf{r}, \tag{1.3}$$

where $\mathbf{r} = \mathbf{n} - \mathbf{M}\boldsymbol{\phi}$ is the vector of timing residuals, obtained from subtracting the input data $\mathbf{n}$ and the timing model evaluated at $\{t_i\}$, here shown as the design matrix $\mathbf{M}$ applied to the vector of timing model parameters $\boldsymbol{\phi}$. The minimization of Equation 1.3 is a generalized least-squares problem, which can be developed further to incorporate a variety of noise terms to model the residuals $\mathbf{r}$ (van Haasteren & Vallisneri 2014; Arzoumanian et al. 2014, 2015a). Such noise modeling is critically important when the signal of interest is itself stochastic and hidden in the noise. In particular, the cutting edge of high-precision pulsar astronomy involves teasing out the gravitational wave signal associated with a background of coalescing supermassive black holes by means of a pulsar timing array.

[10] In the cases of the best MSPs that are free of glitches, timing noise, and ignoring the ISM, these periods of time could be as large as decades.

[11] `http://tempo.sourceforge.net/`

1.4 Pulsar Timing Arrays

Pulsar timing array (PTA) experiments make up one of several attempts to directly detect gravitational waves (GWs). Other ground- and spaced-based efforts aim to measure millihertz and kilohertz frequency gravitational waves by observing the changing path lengths in the arms of laser interferometers. Similarly, Detweiler (1979) first noted that an Earth-pulsar baseline can act as a gravitational wave antenna, which might be sensitive to GWs with periods on the scale of years — the scale over which the pulsar observations are made — corresponding to nanohertz frequencies. The deviations in the pulsar clock arise from the GW-induced Doppler shifts and can be quantified as a limit on the amplitude in gravitational waves from a background of sources or an individual source. Detweiler (1979) concludes with a speculation that the cross-correlation of pulsar signals could identify events extraneous to the pulsars. This idea was carried forward by Hellings & Downs (1983), who first made use of the cross-correlations in pulsar signals to limit the background of GW radiation. Their lasting contribution was to predict that the correlations would be a function of the angular separation of the pulsars, depending on the type of GW source. This allows GWs to be isolated as the culprit for certain signals that are correlated across a network of pulsars.

Among several signal classes, the most promising is thought to be the unresolved, stochastic GW background that is formed from the cosmological history of hierarchical supermassive black hole formation from mergers (Bertotti et al. 1983). For a GW frequency ν, the power-law spectrum of the dimensionless characteristic strain h_c is predicted to be $h_c \propto \nu^{-2/3}$, with an amplitude of $\sim 10^{-15}$ at $\nu = 1$ yr^{-1} (Rajagopal & Romani 1995; Jaffe & Backer 2003). In the power spectrum of the timing residuals, this corresponds to an index of $-13/3$. For $\sim$decade long experiments, this strain amplitude corresponds to RMS timing precisions of at least hundreds of nanoseconds. Obviously, this is achievable only for the best timed MSPs. In fact, only $\sim 20\%$ of field MSPs are "PTA quality". However, as soon as the background GW signal dominates the lowest frequencies of the PTA residuals, the best way to increase sensitivity may be to add as many pulsars as possible to the array, with less regard to their timing quality (Siemens et al. 2013). The other known signal classes include "bright" individual supermassive black hole binaries, cosmic superstrings, and bursts with memory, but the expectations for their detections in the near future are low (Demorest et al. 2013; Arzoumanian et al. 2014, 2015b).

Although the very first realization of a PTA was initiated shortly after the discovery of the first MSPs more than 25 years ago (Foster & Backer 1990), there are currently three PTA experiments working towards opening this brand new window on the Universe. These are the Parkes Pulsar Timing Array (PPTA, Hobbs 2013)[12], which makes use of the Parkes radio telescope, the European Pulsar Timing Array (EPTA, Kramer & Champion 2013)[13], which makes use of the Effelsberg, Lovell, Nançay, Sardinia, and Westerbork radio telescopes, and the North American Nanohertz Observatory for Gravitational Waves (NANOGrav McLaughlin 2013)[14], which will be introduced next. Together, they collaborate as part of the International Pulsar Timing Array[15] (IPTA, McLaughlin 2014), which pools talent, data, and ideas to accelerate the time to a first detection.

[12] `www.atnf.csiro.au/research/pulsar/array/`
[13] `www.epta.eu.org/`
[14] `nanograv.org`
[15] `http://www.ipta4gw.org/`

1.4.1 NANOGrav

The North American Nanohertz Observatory for Gravitational Waves (NANOGrav) is a collaboration between American and Canadian institutions that includes approximately 65 people from ∼16 institutions in a ratio of about 2:1:1 (senior personnel : post-doctoral researchers : graduate students), not including additional staff, and a large number of involved undergraduate and high-school students. Although NANOGrav was officially born around 2007, the first released NANOGrav dataset has its origins in 2005 (Demorest et al. 2013); there are also archival data for a number of NANOGrav pulsars dating back to ∼1998. NANOGrav currently makes ∼monthly timing observations of all its pulsars using the 100-m Robert C. Byrd Green Bank Telescope at the National Radio Astronomy Observatory in Green Bank, West Virginia and the 305-m William E. Gordon Telescope at Arecibo Observatory. NANOGrav is presently experimenting with making timing observations using the Karly G. Jansky Very Large Array, and intends to make augmented use of the soon-to-be-completed 100-meter-class telescope used for the Canadian Hydrogen Intensity Mapping Experiment. The author has been a participating member of NANOGrav since 2009. See Chapter 4 for more details about NANOGrav's observations and the recently finalized 9-year data set (Arzoumanian et al. 2015a).

1.5 The Next Generation of Instrumentation

It is important to note that PTA experiments are currently sensitivity limited. That is, we have not yet reached some fundamental pulsar constraint — e.g. timing noise, pulse jitter, ISM effects — that limits our sensitivity more than the factors that already limit the experiment: collecting area, allocated observing time, receiver sensitivity, bandwidth, number of pulsars, etc. The first four items are encompassed in the radiometer equation, which is the controllable part of the scaling relation for a TOA's uncertainty, $\sigma_{\rm TOA}$ (from Lorimer & Kramer 2005, §8.1.2, Equation 8.2),

$$\sigma_{\rm TOA} \propto \frac{T_{\rm sys}}{A_{\rm eff}\sqrt{t_{\rm obs}\,\Delta f}} \times \frac{P\,\delta^{3/2}}{S_{\rm mean}}. \tag{1.4}$$

$T_{\rm sys}$ is the system temperature, which is related to the receiver sensitivity, $A_{\rm eff}$ is the effective collecting area of the telescope, $t_{\rm obs}$ is the observing time, Δf is the bandwidth, P is the spin period, $S_{\rm mean}$ is the phase-averaged flux density of the pulsar, and $\delta = W/P$ is the duty cycle for a pulse of width W. As is referenced in Chapter 4, several large, blind, and directed pulsar surveys are looking for the small fraction of pulsars that are PTA-quality MSPs; this is the only way to get pulsars with higher $S_{\rm mean}$ and lower δ. With only one or two noteworthy, planned new telescopes, the antennas mentioned above will not see upgrades or replacements, so we can expect $A_{\rm eff}$ to remain about the same. In a related way, $t_{\rm obs}$ is a finite resource to be split among all pulsars and has numerous logistical constraints. Therefore, one guaranteed way to increase a PTA's sensitivity is to develop new instrumentation. Because the improvements in $T_{\rm sys}$ are not expected to be large, the focus has been to increase the instantaneously observed bandwidth Δf, while keeping $T_{\rm sys}$ low. Within Chapters 2 & 4 we provide a background for the movement towards wideband observations and the two generations of instruments used in NANOGrav; for example, in one step, going from $<$ 100 MHz bandwidth to $>$ 600 MHz bandwidth increased the timing precision of many NANOGrav pulsars by factors of ~ 3. The EPTA has recently employed a nearly 3 GHz bandwidth receiver and backend system[16]; the PPTA is following suit, and NANOGrav is also trying to obtain a similar system. All parties

[16] http://www3.mpifr-bonn.mpg.de/staff/pfreire/BEACON.html

involved, however, recognize that traditional timing strategies are quickly outpaced by the growing bandwidths. The issues that need to be addressed motivated much of this dissertation.

1.5.1 The Wide-Bandwidth Problem

Although at first glance a wider bandwidth seems to only provide the benefit of more signal, there are a few important complications to consider. The first issue to realize is that the steep spectra of most MSPs (~ -1.4) limits the practical range of accessible frequencies; a conservative upper cutoff for timing measurements with 100-meter-class telescopes is somewhere between three and five gigahertz. Many PTA MSPs are already almost undetectable at 2.5 GHz. The second issue comes with the frequency dependencies of ISM effects; specifically, pulse broadening from interstellar scattering off of the inhomogeneities in the ISM scales as ν^{-4}. Fortunately for now, many PTA pulsars are relatively nearby and suffer little from the deleterious effects of unmitigated scattering. However, as we continue to add MSPs to the array, it becomes more likely that new MSPs will be less bright and more distant. Also growing at lower frequencies is the background temperature due to Galactic synchrotron radiation, which has a steeper spectrum (~ -2.8) than the average pulsar and yields diminishing returns on going to lower frequencies even if the ISM is not an issue.

Regardless of the above issues, any broadband observations of pulsars will have to contend with intrinsic profile evolution with frequency. For two profiles at disparate frequencies ν_1 and ν_2, one chooses a "fiducial point" in each profile and ascribes this phase to the same longitude on the neutron star surface (Craft 1970). An incorrect alignment is lost in the fit for the DM, which removes a delay $\propto (\nu_1^{-2} - \nu_2^{-2})$. As soon as more than two frequencies are used, however, an incorrect alignment will not get completely absorbed into the absolute DM, but will be propagated as a systematic error in the infinite-frequency TOA(s). This problem of profile alignment was called "the wide-bandwidth problem" in Lommen & Demorest (2013). Part of the issue is alleviated if one has access to a continuous band, where all of the profiles are obtained simultaneously. With the assumption that a ν^{-2} dependence of the delays is valid across the band (and to some degree we know this not to be true (Cordes et al. 2015)), you can constrain the "allowed" alignments by the dispersion law. The full, absolute DM will still be covariant with the choice of profile alignment, however, and it isn't necessarily justified to simply choose the alignment that maximizes the signal-to-noise ratio of the band-averaged profile if there is very significant profile evolution. The changes in DM, which are the important quantities for pulsar timing, will remain unchanged. Disentangling intrinsic profile evolution and dispersive effects has been the subject of much study (e.g., Lommen 2001; Ahuja et al. 2007; Hassall et al. 2012; Liu et al. 2014); we contribute significantly to the resolution of this problem and its associated effects in the remaining chapters.

There are also two other more subtle issues pertaining to wideband observations. For nearby pulsars, diffractive interstellar scintillation acts as a variable weighting function across an observed bandwidth and can introduce timing errors when profile evolution is not modeled (Cordes in prep.). Furthermore, low-weighted frequency channels will produce TOAs with non-Gaussian errors, which must be omitted from the current least-squares timing model fits. A second subtle effect that is now being discussed due to the use of very broad bandwidths and supplementary low-frequency observations is the departure from a pure ν^{-2} law due to a frequency dependent DM. Inhomogeneities in the ISM lead to different paths of propagation for pulses of different frequencies, which necessarily sample slightly different total columns of free electrons. We address both of these concerns as part of the development of our wideband timing protocol.

1.6 Synopsis

- In Chapter 2 we tackle the wide-bandwidth problem directly by developing a novel method for the simultaneous measurement of a TOA and DM. We extend the standard algorithm and make use of arbitrary phase-frequency "portrait" models. In particular, we experiment with modeling pulse portraits as a sum of independently evolving Gaussian components.
  — Pennucci et al. (2014, §1–2)

- In Chapter 3 we make the first test of our methods by examining multi-band radio observations of the bright, fast, high-DM, millisecond pulsar J1824−2452A (M28A); we chose this pulsar particularly because of its dramatic profile evolution and display of scattering. We further validate our measurements through a series of Monte Carlo trials.
  — Pennucci et al. (2014, §3–5)

- In Chapter 4 we make further practical developments of the wideband methodology in order to apply it broadly and carefully to NANOGrav's 9-year profile data set from 37 MSPs. We make timing measurements and model the noise in all pulsars, completing the analysis in parallel with and in comparison to Arzoumanian et al. (2015a).
  — Pennucci et al. in prep

- In Chapter 5 we detail our specialized observational campaign to measure or constrain the Shapiro delay parameters in twelve binary MSPs from the NANOGrav 9-year data set. At least two new neutron star masses may be reported as a result, and an analysis of the Shapiro delay parameters using the wideband data is being completed in parallel with others' efforts.
  — Fonseca/Pennucci et al. in prep

- In Chapter 6 we make multi-band radio and X-ray observations of the Galactic Center magnetar, J1745−2900. Besides determining the absolute phase alignment between the radio and X-ray profiles, we use a simple wideband model of the radio emission to make measurements of the magnetar's DM and scattering parameters as a function of time, as well as reveal its evolving low-frequency radio spectrum.
  — Pennucci et al. (2015)

- In Chapter 7 we quickly summarize these chapters and close with a few remarks.

Chapter 2

Wideband Timing of Radio Pulsars

Note: This chapter comprises the published work: "Elementary Wideband Timing of Radio Pulsars", Pennucci, T. T., Demorest, P. B., & Ransom, S. M. (2014), *The Astrophysical Journal*, 790, 93.

Abstract

We present an algorithm for the simultaneous measurement of a pulse time-of-arrival (TOA) and dispersion measure (DM) from folded wideband pulsar data. We extend the prescription from Taylor (1992) to accommodate a general two-dimensional template "portrait", the alignment of which can be used to measure a pulse phase and DM. We show that there is a dedispersion reference frequency that removes the covariance between these two quantities, and note that the recovered pulse profile scaling amplitudes can provide useful information. We experiment with pulse modeling by using a Gaussian-component scheme that allows for independent component evolution with frequency, a "fiducial component", and the inclusion of scattering. The overall broad application of this new method for dispersion measure tracking with modern large-bandwidth observing systems should improve the timing residuals for pulsar timing array experiments, like the North American Nanohertz Observatory for Gravitational Waves.

2.1 Introduction

The practice of pulsar timing attempts to model the rotation of a neutron star by phase-connecting periodic observations of its pulsed, broadband radio signal. The earliest demonstration of long-term timing observations came relatively soon after the discovery of pulsars (Roberts & Richards 1971). The scientific merits garnered from pulsar timing span astrophysical fields such as planetary science, the interstellar medium, nuclear physics, gravitational wave physics, and are all well-documented (for a review, see eg. Chapter 2, Lorimer & Kramer (2005)).

Pulsar timing and its related experiments have carved out a "sweet spot" in the radio frequency regime that naturally emerges as a trade-off between the pulsar's steep power-law spectrum at the high-frequency end, and the low-frequency drawbacks arising from the pulsed radio signal having to propagate through the ionized interstellar medium (ISM) and the Earth's ionosphere, as well as having to compete with the diffuse background of the galactic synchrotron continuum. The latter has a spectral index in the 1–10 GHz range of ≈ -2.8 (Platania et al. 1998). Population studies have shown that pulsars have an average spectral index around -1.4 at gigahertz frequencies (Bates et al. 2013). The most relevant ISM effect arises from propagation through a homogeneously ionized medium. Interstellar dispersion alters the group-velocity of the radio signal, retarding the arrival of pulses by a time t_{DM} (relative to an infinite frequency signal) according to the cold-plasma dispersion law,

$$t_{\mathrm{DM}} = K \times \mathrm{DM} \times \nu^{-2}, \tag{2.1}$$

where $K \equiv \frac{e^2}{2\pi m_e c} = 4.148808(3) \times 10^3$ MHz^2 cm^3 pc^{-1} sec is called the dispersion constant[1], and DM is the dispersion measure. The dispersion measure is defined as

$$\mathrm{DM} \equiv \int_l n_e \, dl, \tag{2.2}$$

which is the free-electron column density along the path-of-propagation l to the pulsar. The pulse-broadening effect of multi-path propagation through a turbulent, inhomogeneous ISM, known as interstellar scattering, has an even stronger spectral index

[1] K is a combination of the electron charge e, electron mass m_e, and speed of light c. It is common practice in the pulsar community to adopt the approximation $K^{-1} = 2.41 \times 10^{-4}$ MHz^{-2} cm^{-3} pc sec^{-1} (Lorimer & Kramer 2005), which we have used in §3.1 and §3.2.

≈ -4, and becomes increasingly important at lower frequencies for the highest-DM, farthest pulsars (Lorimer & Kramer 2005). Scattering not only broadens the pulsed signal, but delays an intrinsically sharp pulse by an amount roughly proportional to its width, and so is a source of bias in timing measurements. The determination of dispersion measures and effects from scattering have been non-trivial problems concomitant with timing measurements since the beginning (Rankin et al. 1970, 1971).

Nearly all observations taken for (high-precision) pulsar timing experiments are taken within the radio window mentioned above, which lies somewhere in the two decades bounded by about 100 MHz and 10 GHz. The middle decade centered around 1500 MHz seems to be the perennial favorite for timing experiments. Recent developments in pulsar instrumentation and computing over the last 5–10 years have enabled more accurate and sensitive timing measurements. Namely, coherent dedispersion, which completely removes the quadratic time-delay due to a known amount of interstellar dispersion (Hankins & Rickett 1975), required significant advances in computer technology before becoming feasible in real-time on a wide-bandwidth signal.

Historically, observations that implemented coherent dedispersion were limited by computing resources to a bandwidth of order $\sim$100 MHz or less, which is less than most receiver bandwidths. Thus, if one wanted to cover a large portion of the pulsar spectrum, either for timing, spectral, or interstellar medium purposes, several adjacent receiver bands had to be observed separately, which often meant asynchronous measurements and non-contiguous frequency coverage. The implementation of real-time coherent dedispersion to large, instantaneously observed bandwidths has led to the regime wherein the receiver bandwidth (BW) is becoming a limiting factor. The first generation of GHz-bandwidth, coherent dedispersion instruments has been proliferating in the pulsar community for the past several years, beginning with the

Green Bank Ultimate Pulsar Processing Instrument (GUPPI)[2] outfitted for the 100-m Robert C. Byrd Green Bank Telescope (GBT)[3] (DuPlain et al. 2008). GUPPI is an FPGA- and GPU-based system capable of real-time coherent dedispersion of an 800 MHz bandwidth.

The smearing $\delta t_{\rm DM}$ incurred from incorrectly dedispersing a narrow frequency-channel of bandwidth $\Delta\nu = \frac{\rm BW}{n_{chan}}$ and center frequency ν_c by an amount δDM goes as

$$\delta t_{\rm DM} \approx \frac{2K\ \delta{\rm DM}\ \Delta\nu}{\nu_c^3} \approx 8.3\ \left(\frac{\delta{\rm DM}}{\rm cm^{-3}\ pc}\right)\left(\frac{\Delta\nu}{1\ \rm MHz}\right)\left(\frac{\nu_c}{1\ \rm GHz}\right)^{-3}\ \mu{\rm s}. \tag{2.3}$$

This equation demonstrates why it was difficult to observe millisecond pulsars (MSPs) prior to coherent dedispersion; incoherent dedispersion shifts subbands of the data based on the assumed DM, meaning δDM was equivalent to the full, true DM. Therefore, $\delta t_{\rm DM}$ could easily exceed the pulse period ($P_s \lesssim 10$ ms for MSPs). Equation 2.3 also highlights why instantaneous measurements of the DM are necessary when one has access to data from a large fractional bandwidth; the delay across an 800 MHz bandwidth at $\nu_c = 1500$ MHz from incorrectly dedispersing by 0.01 cm^{-3} pc is ~20 μs, which is comparable to the width of sharp features observed in the average pulse profiles of "high-precision" MSPs (Jacoby et al. 2003).

In particular, tracking the dispersion measure changes in MSP observations is necessary for mitigating the timing residuals used in gravitational wave searches with a pulsar timing array (PTA) (You et al. 2007). As part of the Parkes Pulsar Timing Array[4] project (Hobbs 2013), Keith et al. (2013) developed a method to correct for inaccurate dispersion measures based on modeling the multi-frequency timing residuals. However, the authors also postulate that more accurate DM variations could be

[2]`www.safe.nrao.edu/wiki/bin/view/CICADA/NGNPP`

[3]The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

[4]`www.atnf.csiro.au/research/pulsar/array/`

measured from wideband receivers, which ameliorate the difficulties of aligning pulsar data taken with different receivers in different epochs.

The desire for very broadband pulsar observations (i.e. with significantly high fractional bandwidths, $\gtrsim$1) necessitates new, unique receiver designs that can cover much of the frequency range once concatenated from disjoint observations. Wideband receivers and their complimentary, real-time coherent dedispersion backends will quickly facilitate developments in all realms of pulsar astrophysics, including studies of the pulsar spectrum, magnetosphere, and ISM properties. One such instrument, called the Ultra-Broad-Band (UBB) receiver and associated backend[5], has been recently installed at the Effelsberg 100-m Telescope and covers a frequency range from $\sim$600 – $\sim$3000 MHz.

However, the current method for making pulse time-of-arrival (TOA) measurements that is used almost ubiquitously in the pulsar timing community does not use all of the information contained in new broadband observations. In summary, the protocol employs frequency-averaged pulse profiles as models of the pulsar's signal for entire receiver bands, which ignores any profile evolution intrinsic to the pulsar or imposed by the ISM. Both intrinsic profile evolution and DM changes are usually taken into account in the timing model for the pulsar's rotation, but there is no modeling of the effects from scattering or scintillation. Arbitrary phase-offsets (known as "JUMPs") are introduced to align disparate template profiles that are used to measure TOAs from different frequency bands. Multi-channel TOAs are also parameterized by both a quadratic delay (proportional to the DM) and an arbitrary function to remove residual frequency structure from otherwise unmodeled profile evolution.

These methods are ad hoc and incomplete in that they were developed as the availability of bandwidth and multi-frequency observations became a "problem" (cf.

[5] `www3.mpifr-bonn.mpg.de/staff/pfreire/BEACON.html`

"the wide-bandwidth problem" (Lommen & Demorest 2013)), and were appropriate when observations covered a narrow bandwidth: phase JUMPs account for profile evolution occurring in frequency gaps that are not observed. It seems natural in the era of wideband receivers — when frequency evolution *is* observed in the band — to devise a method for TOA measurement that includes a frequency-dependent model of the average pulse profile. In doing so, it becomes straightforward to include a simultaneous measurement of the dispersion measure. As we will show, a very simple extension to the algorithm that is currently used is a first step in a more comprehensive and necessary description of the received pulsar signal.

2.2 The Algorithm

2.2.1 Background

We assume that the recorded pulsar signal is cyclostationary for a given frequency, meaning the observed time-series data can be coherently folded modulo a pre-existing timing model to obtain an average signal shape that is stable with time. This time-integrated light curve is often called a "pulse profile", which we label as $D(\varphi)$. The quantity $\varphi = \varphi(t_{obs})$ represents the rotational phase of the neutron star at a particular moment in time t_{obs}, which is recorded by an observatory clock and later transformed into a more useful temporal coordinate system.

The central step in determining a pulse time-of-arrival is to measure the relative phase shift $\phi \in$ [-0.5, 0.5) between $D(\varphi)$ and a standard template profile, $P(\varphi)$, which is supposed to represent the noise-free average of the intrinsic pulse profile shape at the observed frequency. In practice, the signal has been discretely and evenly sampled so that $D(\varphi)$ becomes $D(\varphi_j = (j + 0.5)/n_{bin}) \equiv D_j$, where j runs from 0 to $n_{bin} - 1$, and n_{bin} is the number of phase bins in the profile. For pulsar timing purposes, the sampling time (and therefore the number of bins in the profile) is chosen to be appropriately small so that all meaningful information about the pulse profile with respect to the noise level is preserved in D_j.

A simple way to obtain a lag between D_j and a template profile P_j is to interpolate a maximum point in the discrete time-domain cross-correlation of the two functions. Appendix A of Taylor (1992) prescribes a Fourier frequency-domain technique for measuring the phase shift that has been used virtually ubiquitously for the past two decades in the pulsar timing community. Besides the computational simplicity that is a consequence of the Fourier cross-correlation theorem, the reason for this ubiquity is because frequency domain techniques give very precise, accurate shifts for low duty-cycle pulsars that correspond to less than a single time bin (Taylor 1992; Hotan et al. 2005). Colloquially, this routine came to be known as FFTFIT, which is the designation we will use hereafter. The advantage of FFTFIT is that a finite number of continuously-valued Fourier phases (instead of discrete time lags) are combined to interpolate a precise phase measurement. An alternate formulation of FFTFIT can be found in Chapter 2 of Demorest (2007), which also recognizes that FFTFIT amounts to a cross-correlation completed in the frequency domain. We have drawn from Demorest (2007) as a starting point for the mathematical framework, and have borrowed some of its notation in what follows.

2.2.2 Description

Because we are concerned with measurements of a wideband pulsar signal, we describe the observed pulse profile also as a function of frequency ν, which we denote by $D(\nu, \varphi)$, and refer to as a "pulse portrait". Similarly, the template portrait is

$P(\nu, \varphi)$ and a simple model for the observed data is

$$D(\nu, \varphi) = B(\nu) + a(\nu)P(\nu, \varphi - \phi(\nu)) + N(\nu, \varphi), \tag{2.4}$$

where $\phi(\nu)$ will contain information about chromatic and achromatic phase shifts, $B(\nu)$ is effectively the bandpass shape of the receiver (analogously, B is the "DC" or "bias" term when considering only a single frequency, as in FFTFIT), $a(\nu)$ is a multiplicative scale factor that can represent scintillation, and $N(\nu, \varphi)$ is additive noise. $N(\nu, \varphi)$ is often assumed to be stationary and normally distributed with variance $\sigma^2(\nu)$, so that $N(\nu) \sim \text{Normal}(0, \sigma^2(\nu))$. In the absence of radio-frequency interference (RFI), the noise in most pulse profiles is radiometer-noise dominated, which is highly Gaussian. There are numerous methods for the removal of the bandpass shape $B(\nu)$ (which can be thought of as the frequency-dependent mean of the noise term $N(\nu)$), or one could follow an analogous treatment of the bias term in Taylor (1992). One simple solution is to start all of the Fourier phase sums in the below equations at $k = 1$, as we have done for our implementation.

Again, in practice the signal is discretized into n_{bin} phase bins, but also into n_{chan} frequency channels with center frequencies ν_n. We index each of the above frequency-dependent quantities with the letter n (eg. ϕ_n, D_{nj}, P_{nj}). The question of determining n_{chan} will be revisited in §3.2. The Discrete Fourier Transform (DFT) is a linear transformation, so taking a one-dimensional DFT of Equation 2.4 with respect to rotational phase φ, and making use of the discrete Fourier shift-theorem (Bracewell 2000) implies

$$d_{nk} = a_n p_{nk} e^{-2\pi i k \phi_n} + n_{nk} \qquad (k > 0), \tag{2.5}$$

where $i = \sqrt{-1}$, k indexes the Fourier frequencies, and the DFT of a series F_j is

$$f_k = \sum_{j=0}^{n_{bin}-1} F_j e^{-2\pi ijk/n_{bin}}. \tag{2.6}$$

The primary quantities of interest ϕ_n, and the scaling parameters a_n in Equation 2.5 can be found by minimizing the sum of the squares of the residuals between the data d_{nk} and the shifted, scaled template p_{nk}, weighted by the noise[6] in each frequency channel $\sigma_n'^2$. In other words, we seek to minimize the statistic

$$\chi^2(\phi_n, a_n) = \sum_{n,k} \frac{|d_{nk} - a_n p_{nk} e^{-2\pi ik\phi_n}|^2}{\sigma_n'^2}. \tag{2.7}$$

It is useful at this point to make note of the fact that at a given frequency ν_n, the above expression is equivalent to the FFTFIT prescription. The fundamental difference in this approach, besides allowing for an arbitrary evolution of the pulse profile with frequency encoded in p_{nk}, is that we perform a global fit for both an achromatic phase ϕ°_{ref} and a dispersion measure DM by implementing the constraint

$$\phi_n = \phi^\circ_{ref} + \frac{K \times \mathrm{DM}}{P_s}\left(\nu_n^{-2} - \nu_{ref}^{-2}\right), \tag{2.8}$$

where P_s is the spin period and ν_{ref} is the dedispersion reference frequency. This constraint reduces our minimization problem from having $2n_{chan}$ parameters, to $n_{chan}+2$ (i.e. $\chi^2(\phi_n, a_n) \rightarrow \chi^2(\phi^\circ_{ref}, \mathrm{DM}, a_n)$). Ideally, we want to know what ϕ°_{ref} is for $\nu_{ref} = \infty$. However, we have chosen the above parameterization for the delays in each frequency channel, as opposed to the more specific infinite-frequency case of Equation 2.1, because it allows us to find a reference frequency that gives zero co-

[6]Assuming Gaussian noise, the noise variance $\sigma_n'^2$ in each frequency channel of d_{nk} is greater than σ_n^2 in D_{nj} by the factor $n_{bin}/2$.

variance between the estimates of the phase and dispersion measure. The form of the covariance between the estimates of ϕ°_{ref} and DM is given in §2.3, which recommends that we choose ν_{ref} wisely (see also §3.2).

By following a similar procedure as that written in Demorest (2007), and expanding and simplifying Equation 2.7 we obtain

$$\chi^2(\phi^{\circ}_{ref}, \mathrm{DM}, a_n) = S_d + \sum_n a_n^2 S_{p,n} - 2\sum_n a_n C_{dp,n} \tag{2.9}$$

where we have made use of the definitions

$$S_d \equiv \sum_{n,k} \frac{|d_{nk}|^2}{\sigma'^2_n}, \tag{2.10a}$$

$$S_{p,n} \equiv \sum_{k} \frac{|p_{nk}|^2}{\sigma'^2_n}, \tag{2.10b}$$

and

$$C_{dp,n}(\phi_n) \equiv \Re\Big\{\sum_k \frac{d_{nk}p^*_{nk}e^{2\pi ik\phi_n}}{\sigma'^2_n}\Big\}. \tag{2.10c}$$

The first two definitions are functions solely of the data and the model portraits. If one considers discrete values of ϕ_n for a particular frequency channel n ($\phi_{nj} = j/n_{bin}$), the third definition contains the inverse DFT of a multiplication of the data and the model, which is the same as the discrete cross-correlation of the time-domain quantities D_{nj} and P_{nj}. This definition highlights the fact that both FFTFIT and our extension of it across a discretized bandwidth can be thought of as cross-correlation techniques.

We can further simplify our minimization problem by recognizing that at the global minimum of the χ^2 expression in Equation 2.9, all of the first derivatives vanish. Therefore, we only need to seek out a minimum of Equation 2.9 in the subspace where its partial derivatives with respect to all of the a_n parameters are zero. Solving for these a_n as a function of the other parameters leads to the constraint

$$a_n = \frac{C_{dp,n}}{S_{p,n}}, \tag{2.11}$$

which is inserted in Equation 2.9 to reduce our minimization problem to a two-parameter function,

$$\chi^2(\phi^\circ_{ref}, \mathrm{DM}) = S_d - \sum_n \frac{C^2_{dp,n}}{S_{p,n}}. \tag{2.12}$$

We retain the use of the label χ^2 to emphasize that the above function is a subspace of Equation 2.9, and shares the global minimum that we seek. In practice, one needs to maximize only the strictly positive second term in the above equation, since it contains all of the phase and dispersion information, and the first term is a constant function of the data. It is easy to see that, for negligible profile evolution, if the dispersion measure is zero or, equivalently, the data have been correctly dedispersed for that observation's epoch, then this algorithm is akin to averaging TOAs obtained in the usual way using individually aligned templates. However, if the pulsar's DM needs to be measured at every epoch, and profile evolution should be accounted for — which are both likely true for most observations of MSPs with wideband receiver systems — we claim that this is a natural extension to how TOAs are currently procured. At the very least, this algorithm should perform no worse than traditional techniques.

We derive errors and covariances for the maximum-likelihood estimates of the parameters in §2.3, but here we wish to underscore that it is possible to analytically determine a dedispersion reference frequency ν_{zero} that yields zero covariance between the estimates for ϕ°_{ref} and DM, which is tested in §3.2. Lastly, we note that Liu et al.

(2014) have contemporaneously developed a very similar frequency-dependent TOA algorithm independent of our efforts, which may be employed as part of the European Pulsar Timing Array[7] project (Kramer & Champion 2013).

2.2.3 Implementation

Software

We have implemented our wideband timing algorithm in publicly-available `python` code[8], which also includes a Gaussian-component-based portrait modeling routine, which is described below. The code utilizes the `python` interface to the pulsar data analysis package `PSRCHIVE`[9] (Hotan et al. 2004; van Straten et al. 2012), as well as recent versions of `numpy`[10], the optimization functions in `scipy`[11], and the non-linear least-squares minimization package `lmfit`[12].

The minimization of the function in Equation 2.12 is performed by a truncated Newton algorithm that comes packaged in `scipy`. The initial phase parameter value is estimated by using a one-dimensional brute-force routine in `scipy`, which is performed on the frequency-averaged data and template. In order to do this, the data are dedispersed with respect to an estimate for ν_{zero}, which can only be determined after the minimum is found. The nominal DM from the pulsar's ephemeris is used in this dedispersion and also as the initial DM parameter value in the global fit. It is also possible to exclude fitting for a DM and only determine a phase. The default behavior in the code transforms the best-fit phase estimate $\hat{\phi}^{\circ}_{ref}$ to reference ν_{zero}, which gives

[7] `www.epta.eu.org/`
[8] `www.github.com/pennucci/PulsePortraiture`
[9] `www.psrchive.sourceforge.net/`
[10] `www.numpy.org/`
[11] `www.scipy.org/`
[12] `www.newville.github.io/lmfit-py/`; `lmfit` is a Levenberg-Marquardt algorithm that we use for the modeling code.

the smallest and uncorrelated errors for the TOA and DM (see §2.3).

It is important to use barycentric frequencies for ν_n, otherwise the Earth's orbit induces an apparent yearly oscillation of the DM from the Doppler-shifted frequencies. Alternatively, one can simply propagate the Doppler factor Γ through the frequency and temporal terms of Equation 2.1 to correct the observed "topocentric" dispersion measure DM_{topo},

$$\mathrm{DM} = \frac{\mathrm{DM}_{topo}}{\Gamma}, \tag{2.13}$$

where

$$\Gamma \equiv \sqrt{\frac{1+\beta}{1-\beta}}, \tag{2.14}$$

$$\beta \equiv \frac{v}{c}, \tag{2.15}$$

and v, the projected velocity of the observatory onto the line-of-sight, is positive for growing separation. With respect to the demonstration in the next section, our source is close to the ecliptic plane and has a large DM, so this correction was essential.

Portrait Modeling

It is obvious that there is freedom in the choice of model portrait to use and we stress that any arbitrary model can be used in the above algorithm for phase and DM measurements. We have experimented almost exclusively with analytic Gaussian-component models, but it is also feasible to find an interpolation scheme based on an average of all the data portraits (or, for example, a principal component analysis approach). However, Gaussian-component modeling has been used extensively in the literature (for instance, see Foster et al. (1991); Kramer et al. (1998); Lommen (2001); Ahuja et al. (2007); Hassall et al. (2012)) and is a simple way to generate analytic noise-free templates. We model pulse profile evolution with independently changing

Gaussian components g_i of the form

$$P(\nu, \varphi) = \sum_i g_i(\nu, \varphi), \tag{2.16}$$

where

$$g_i(\nu, \varphi | A_i, \varphi_i, \sigma_i) = A_i(\nu)\exp\Big(- 4\ln(2)\frac{(\varphi - \varphi_i(\nu))^2}{\sigma_i(\nu)^2}\Big). \tag{2.17}$$

We choose to model the positions φ_i, widths (FWHM) σ_i, and amplitudes A_i as power-law functions of frequency,

$$X_i(\nu | X_{\circ,i}, \alpha_{X,i}, \nu_\circ) = X_{\circ,i}\Big(\frac{\nu}{\nu_\circ}\Big)^{\alpha_{X,i}}, \tag{2.18}$$

for Gaussian parameter X and model reference frequency $\nu_\circ$. We also include linear functions for φ_i and σ_i in the code. The modeling code allows flexibility for any of the parameters to be fixed; for example, a "fiducial component" with no positional change as a function of frequency can be selected. Most MSP portraits we have experimented on seem to be sufficiently characterized by a few to roughly a dozen or so Gaussian components. We also include an option to include scattering in the fit for the model via a convolution with a one-sided exponential,

$$P(\nu, \varphi) = P_{unscattered}(\nu, \varphi) * e^{-\frac{\varphi P_s}{\tau(\nu)}} H(\varphi) \tag{2.19}$$

where

$$\tau(\nu) = \tau_\circ\Big(\frac{\nu}{\nu_\circ}\Big)^{\alpha_{scat}}, \tag{2.20}$$

H is the Heaviside step function, and we have assumed $\alpha_{scat} = -4.0$ (Bhat et al. 2004). One could imagine extending our algorithm to include a variable scattering parameter in the fit to the data, instead of fixing it in the model. The benefits, applicability, and practical limitations of doing this are currently being investigated by the authors. The details of pulse portrait modeling and its physical interpretations are beyond the scope of this paper, but we demonstrate one application of Gaussian modeling in the next section.

Finally, one subtlety that we did not address is the averaging of the model within each frequency channel to match the channel bandwidth of the data. Insofar that the aim is to have a greater number of channels, this should have a negligible effect. Presumably, each channel's profile evolution is minute and, as given in Equation 2.3, the channel smearing from an inaccurate DM is also small. Similarly, we assumed that a sufficient number of phase bins are used so that all of the harmonic content of the profile is retained and an averaging of the model into phase bins is well-approximated by the Gaussian model values. However, a more rigorous representation for P_{nj} would be one that multiplies the model by a sampling function that has been convolved with both the channel-width and bin-size of the data.

2.3 Appendix: Covariance & Error Estimates of Fitted Parameters

We will denote our least-squares estimate of the parameters $\theta = \{\phi^{\circ}_{ref}, \mathrm{DM}, a_n\}$ by $\hat{\theta}$. Equation 2.9 has a minimum at $\hat{\theta} = \{\hat{\phi}^{\circ}_{ref}, \hat{\mathrm{DM}}, \hat{a}_n\}$, where $\hat{a}_n = \frac{\hat{C}_{dp,n}}{S_n} = \frac{C_{dp,n}(\hat{\phi}^{\circ}_{ref}, \hat{\mathrm{DM}})}{S_n}$ (cf. Equation 2.11). We can approximate the $(n_{chan}+2) \times (n_{chan}+2)$ covariance matrix of the parameters by the inverse of the curvature matrix $\boldsymbol{\kappa}$, which can be derived from a Taylor expansion of the χ^2 function near its minimum. The entries of the curvature matrix at the minimum point are given by

$$\kappa_{kl} = \frac{1}{2}\frac{\partial^2\chi^2(\theta)}{\partial\theta_k\partial\theta_l}\bigg|_{\hat{\theta}}, \tag{2.21}$$

which is one-half of the Hessian. However, because we are primarily interested in only the errors and covariances of $\hat{\phi}^{\circ}_{ref}$ and $\hat{\mathrm{DM}}$, we will only calculate here the terms of the 2×2 Hessian for an arbitrary point of the function $\chi^2(\phi^{\circ}_{ref}, \mathrm{DM})$ given in Equation 2.12. Inverting this matrix to arrive at the 2×2 covariance matrix of interest is trivial, particularly because there is a reference frequency ν_{zero} that gives zero covariance between $\hat{\phi}^{\circ}_{ref}$ and $\hat{\mathrm{DM}}$[13]. One can arrive at the same results for the corresponding entries of the "full" covariance matrix by inverting the matrix in Equation 2.21 and inserting the values $a_n = \frac{C_{dp,n}}{S_n}$. The three unique second-derivatives of Equation 2.12 are

$$\frac{\partial^2\chi^2}{\partial\phi^{\circ 2}_{ref}} = -2\sum_n w_n, \tag{2.22a}$$

$$\frac{\partial^2\chi^2}{\partial\mathrm{DM}^2} = -2\sum_n w_n\left[\frac{K}{P_s}(\nu_n^{-2} - \nu_{ref}^{-2})\right]^2, \tag{2.22b}$$

and

$$\frac{\partial^2\chi^2}{\partial\phi^{\circ}_{ref}\partial\mathrm{DM}} = -2\sum_n w_n\left[\frac{K}{P_s}(\nu_n^{-2} - \nu_{ref}^{-2})\right], \tag{2.22c}$$

where

$$w_n \equiv S_{p,n}^{-1} \times (C'^2_{dp,n} + C''_{dp,n}C_{dp,n}), \tag{2.23}$$

and the derivatives of $C_{dp,n}$ are with respect to ϕ_n. Requiring that the cross-term in Equation 2.22c be equal to zero leads us to the zero-covariance dedispersion reference

[13]To clarify, the covariances with the $\hat{a}_n$ estimates are already included in the 2×2 covariance matrix; there is only zero covariance at ν_{zero} between the fitted phase and DM.

frequency ν_{zero},

$$\nu_{zero} = \sqrt{\frac{\sum_n w_n}{\sum_n w_n \nu_n^{-2}}}. \tag{2.24}$$

Because Equation 2.10c is a function of ϕ_n, we can transform the least-squares estimate $\hat{\phi}^{\circ}_{ref}$ to an estimate that has zero covariance with $\hat{\mathrm{DM}}$ and ensure we are still at the minimum point,

$$\hat{\phi}^{\circ}_{zero} = \hat{\phi}^{\circ}_{ref} + \left[\frac{K \times \hat{\mathrm{DM}}}{P_s}\left(\hat{\nu}_{zero}^{-2} - \nu_{ref}^{-2}\right)\right]. \tag{2.25}$$

Under this transformation of the ϕ°_{ref} coordinate, the Hessian is diagonal and the variances of the estimates $\hat{\phi}^{\circ}_{zero}$ and $\hat{\mathrm{DM}}$ are simply twice the inverse of Equations 2.22a and 2.22b, respectively. A derivation starting with the full $(n_{chan}+2)\times(n_{chan}+2)$ Hessian confirms that these errors incorporate the covariances with the a_n. Furthermore, the Monte Carlo results from §3.2 and Figure 3.10 show that we have accurately been able to calculate covariances down to low SNR levels. The default output TOAs from our `python` code references the phase estimates and TOAs to ν_{zero}, which is our recommendation for any similar implementation.

Chapter 3

Demonstration on PSR J1824–2452A & Monte Carlo Tests

Note: This chapter comprises the published work: "Elementary Wideband Timing of Radio Pulsars", Pennucci, T. T., Demorest, P. B., & Ransom, S. M. (2014), *The Astrophysical Journal*, 790, 93.

Abstract

We showcase the algorithm from the previous chapter using our publicly available code on three years of wideband data from the bright millisecond pulsar J1824−2452A (M28A) from the Green Bank Telescope, and a suite of Monte Carlo analyses validates the algorithm. By using a simple model portrait of M28A we obtain DM trends comparable to those measured by standard methods, with improved TOA and DM precisions by factors of a few. Therefore, we expect that the measurements from our algorithm will yield precisions at least as good as those from traditional techniques, but is prone to fewer systematic effects and is without ad hoc parameters.

3.1 Demonstration with MSP J1824−2452A

3.1.1 M28A Dataset and Model Portrait

Pulsar J1824−2452A (M28A, hereafter) is a highly energetic, bright, isolated 3.05 ms pulsar in the globular cluster M28 (Lyne et al. 1987; Johnson et al. 2013). We chose this MSP as a demonstrative case-study because it has a large dispersion measure ($\approx$120 cm^{-3} pc), a large DM gradient (several $\times 10^{-3}$ cm^{-3} pc yr^{-1}) (Backer et al. 1993; Cognard & Lestrade 1997; Keith et al. 2013), a complex profile with broad and narrow features, and because it shows component evolution across the frequency range 720 – 2400 MHz (Foster et al. 1991).

The M28A dataset presented here consists of 25 epochs of multi-frequency observations spanning more than three years from the Green Bank Telescope. The data were obtained with GUPPI beginning in February 2010 soon after the implementation of its real-time coherent dedispersion capability, which was utilized for the taking of these observations in search-mode (i.e as unfolded time-series). Each of the time-series for the 512 channels across each frequency band were dedispersed at a nominal average DM for the globular cluster, 120 cm^{-3} pc, and then folded using a predetermined ephemeris for M28A. The native resolution of the data is 10.24 μs, which is sufficient to resolve the profile, although we folded the data at nearly twice this resolution, resulting in 512 phase bins. A more technical description of these data and their calibration is provided in Bilous et al. (2015). Table 3.1 summarizes the epochs of the observations presented here.

Figure 3.1 shows a concatenated portrait of several epochs of the M28A data, displaying an effective bandwidth of $\sim$1.5 GHz. The complexity of the portrait is evinced by its asymmetries, its non-Gaussian features, the exchanging dominance of components from differing spectral indices, and the presence of an obvious scattering tail at the lower frequencies. To make this portrait, five high signal-to-noise ratio (SNR) epochs were selected from each set of 1500 MHz and 2000 MHz observations, they were each averaged together based on the ephemeris, and then joined in tandem along with the 820 MHz observation in a fit for the two-dimensional Gaussian model, as described in §2.2.3. The fit included nuisance phase and DM parameters for each band, as well as a scattering timescale $\tau_{\circ}$. In effect, the nuisance parameters attempted to "align" the data so that the Gaussian parameters can be optimized.

Following the suggestion in Foster et al. (1991), we modeled the widths of M28A's components with power-law functions, and had less success when trying linear functions. To obtain initial parameters for the two-dimensional model, we fit ten Gaussian components to a profile referenced at 1500 MHz, representing 200 MHz of bandwidth averaged. We chose the dominant component at 1500 MHz to be a fixed "fiducial component". The parameters of the fitted model are given in Table 3.2. MSPs like M28A exemplify how the choice of a "fiducial point" is not simple (eg. see Craft

Table 3.1. J1824−2452A: Summary of GBT Observations and DMs

Epoch [UTC]	MJD [day]	ν_c [MHz]	Length [min]	ΔDM [$\times 10^{-3}$ cm^{-3} pc]
2010-02-11	55238.72	1500	43.3	-2.4 ± 0.2
2010-05-20*	55336.35	2000	129.0	-2.2 ± 0.4
2010-08-11	55419.15	2000	166.3	-0.6 ± 0.4
2010-10-05*†	55474.00	820	159.4	0.13 ± 0.07
2010-10-20*	55489.93	1500	149.3	0.18 ± 0.06
2011-03-05	55625.58	1500	157.2	2.58 ± 0.08
2011-04-04	55655.55	1500	98.8	1.56 ± 0.09
2011-04-13*	55664.42	1500	154.2	1.42 ± 0.06
2011-07-02*	55744.25	1500	149.3	-0.12 ± 0.06
2011-09-29	55833.98	1500	154.4	0.28 ± 0.06
2012-01-06	55932.72	1500	145.3	-1.91 ± 0.06
2012-04-09*	56026.45	1500	149.3	-2.81 ± 0.04
2012-04-15*	56032.42	2000	161.4	-2.8 ± 0.2
2012-07-03	56111.25	1500	138.3	-2.23 ± 0.06
2012-10-07†	56207.96	1500	152.3	-0.55 ± 0.06
2013-01-06	56298.70	1500	149.3	1.54 ± 0.06
2013-04-08*	56390.48	1500	176.4	0.02 ± 0.05
2013-04-15*†	56397.46	2000	164.4	0.1 ± 0.3
2013-05-06	56418.29	2000	82.2	1.0 ± 0.5
2013-05-09	56421.44	2000	83.2	1.5 ± 0.7
2013-05-11*	56423.40	2000	77.2	0.8 ± 0.5
2013-05-13*	56425.43	2000	82.2	1.4 ± 0.5
2013-05-18	56430.27	2000	83.2	0.7 ± 0.5
2013-05-24	56436.43	2000	59.1	1.6 ± 0.7
2013-05-31	56443.18	2000	80.2	0.6 ± 0.4

Note. — The columns are the UTC YYYY-MM-DD observation date, the Modified Julian Date, the center frequency, the total integration time, and the measured DM with 1σ errors. The DMs had the nominal (unweighted) average value of 119.88818 cm^{-3} pc subtracted. There are 11 epochs observed at 2000 MHz (800 MHz BW), 13 at 1500 MHz (800 MHz BW), and 1 at 820 MHz (200 MHz BW). The fractional bandwidths are approximately 0.25, 0.53, and 0.40 for the 820, 1500, and 2000 MHz data, respectively. The starred epochs were used in the fit for the Gaussian model, and epochs with a dagger are shown as part of Figure 3.2.

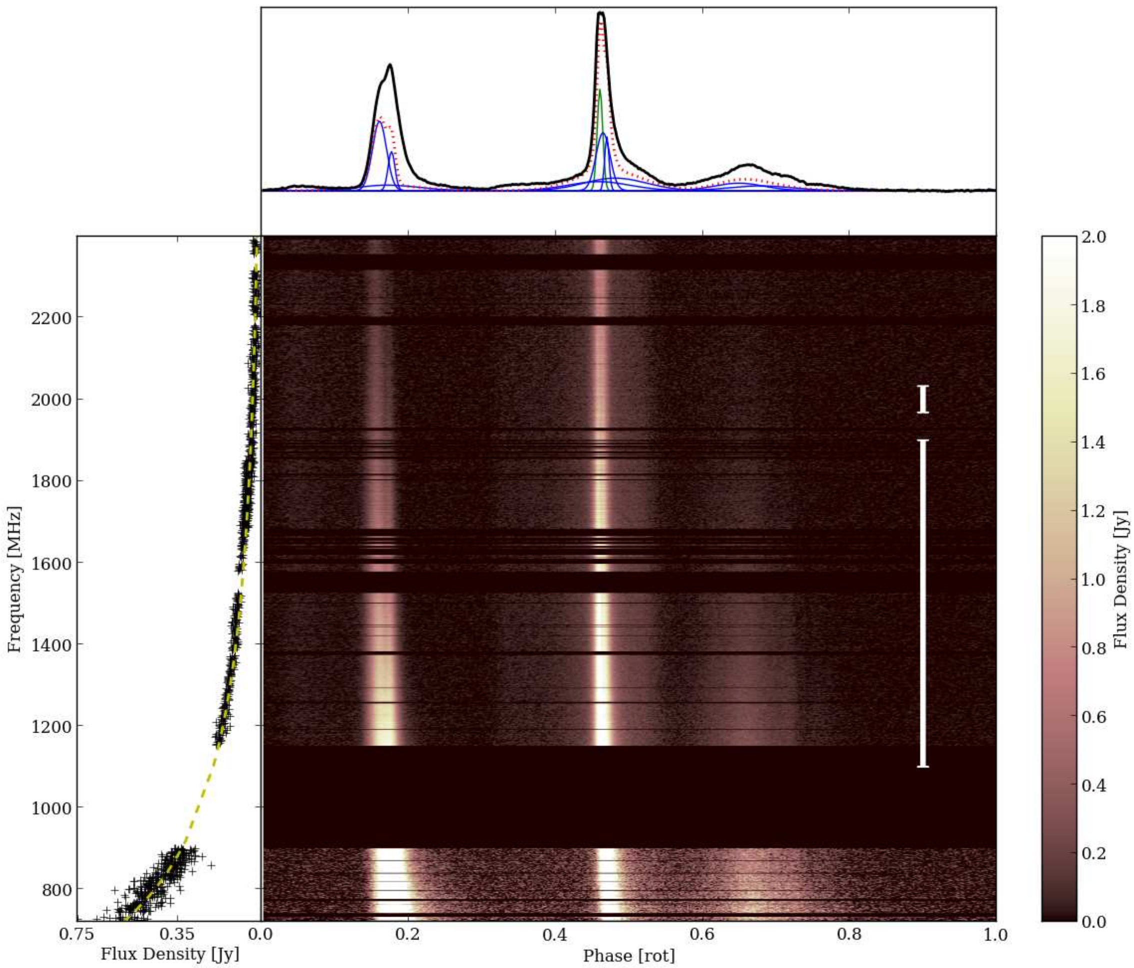

Fig. 3.1 – A portrait of concatenated M28A data from time-averaged observations taken at 820 MHz, 1500 MHz, and 2000 MHz. See Table 3.1 for the epochs used in the averages and the text for a complete description. The horizontal gaps are where radio-frequency interference was excised, but the widest one is a coverage gap between receivers. The two white vertical bars show the bandwidth coverage offered by previous coherent dedispersion backends (64 MHz) and current ones like GUPPI (800 MHz). Note that the UBB has an instantaneous bandwidth of almost 1.5 times that shown above.

(1970)) because the profile has no obvious symmetries, and the dominant component changes as a function of frequency. The model and the residuals are shown in Figure 3.3.

The thick solid black line in the top panel of Figure 3.1 represents the frequency-averaged light curve of the aligned data. This profile marginalizes over all of the frequency structure and scattering tails, and so it would be imprudent to use such a profile as a template for obtaining phase measurements. To get a sense of the model for this data, the thinner blue curves show the Gaussian components from the model fit at 1500 MHz, and the tallest (green) curve is the "fiducial component". The components are shown unscattered and scaled (relative to the black light curve) for clarity; the red dotted line is the sum of the components, including scattering. The phase-averaged spectral flux density profile in the left panel was fit with a power-law (yellow dashed line). We obtained a spectral index of -2.36 ± 0.02, although it appears as though the flux is not perfectly modeled by a single power-law. Details of M28A's spectra from this dataset, including an analysis of its polarization properties, are also presented in Bilous et al. (2015).

Some of the subtle profile evolution for this pulsar can be seen in the top panel of Figure 3.2, which consists of timing residuals as a function of frequency (see §3.1.2). As is evident from the top panel, the use of an average template profile to measure TOAs for a band, or portion thereof, will produce a different residual as a function of frequency based on the profile's departure from the frequency-averaged template. If this frequency-dependent bias were constant, in would be absorbed into the timing model, but varying scintillation patterns can change which segments of the bias are weighted more significantly (or, similarly, what the frequency-averaged profile looks like), thereby introducing random systematic noise into the timing residuals.

Table 3.2. J1824−2452A: Gaussian Model Parameters

i	$\varphi_{\circ,i}$ [rot]	$\alpha_{\varphi,i}$	$\sigma_{\circ,i}$ [% rot]	$\alpha_{\sigma,i}$	$A_{\circ,i}$	$\alpha_{A,i}$
1	-0.00180	-0.693	10.00	0.3	0.09	-1.2
2	0.00000	0.000	0.88	0.2	1.11	-1.9
3	0.00410	-0.021	2.24	0.3	0.63	-2.1
4	0.00932	-0.017	0.69	-0.1	0.58	-1.4
5	0.02078	0.280	9.96	-2.0	0.13	-0.1
6	0.18894	-0.006	7.93	0.4	0.08	-3.6
7	0.21877	-0.124	10.00	0.0	0.05	-3.3
8	0.70012	-0.007	2.24	-0.1	0.75	-3.0
9	0.71061	-0.025	9.98	5.3	0.05	-6.1
10	0.71651	-0.001	1.09	0.2	0.42	-3.5
$\tau_\circ$	4.57 μs					
$\nu_\circ$	1500.00 MHz					

Note. — The column headers are defined in Equation 2.18. The components are ordered by phase; Figures 3.1, 3.3, and 3.4 have been rotated for clarity. The second component listed is the "fiducial component". A limit of 0.1 rotations was placed on the FWHM width of the components to prevent runaway for small-amplitude components. The precision of all the parameters is arbitrary, since we offer no interpretation of the model in this paper. The reference frequency for the toy model is 1500 MHz and a scattering kernel corresponding to a fitted scattering timescale of $\tau_\circ \approx 5$ μs at 1500 MHz was applied to the model (cf. Equations 2.19 and 2.20). The point estimate of the scattering timescale is marginally consistent with that found independently from a separate analysis of giant pulses in this M28A data (Bilous, private communication).

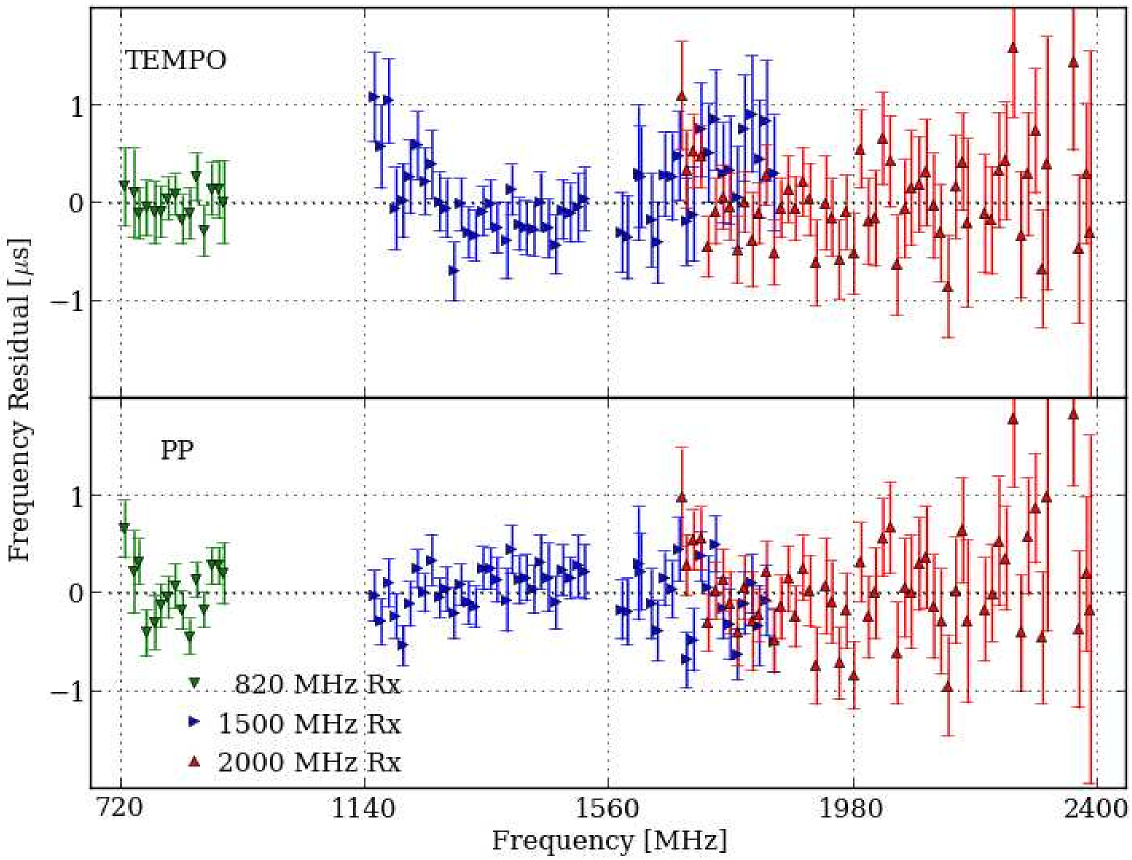

Fig. 3.2 – Timing residual structure demonstrating profile evolution as a function of frequency for some M28A data after the DMs have been measured (see §3.1.2 for discussion). Table 3.1 denotes which epochs are shown. Each point represents 12.5 MHz of bandwidth averaged. The top panel ("TEMPO") does not account for any profile evolution, but only uses a single template profile per receiver band; the opposing trends in the overlapping region between the 1500 and 2000 MHz residuals signify that there is no continuous frequency-dependent model. A comprehensive, global model for the evolution would ideally show flat residuals. The "residuals" from applying our algorithm with the Gaussian toy-model are shown in the lower panel ("PP"), showing the best results in the 1500 MHz data. Next generation wideband receivers will simultaneously cover more than this entire spectrum at once.

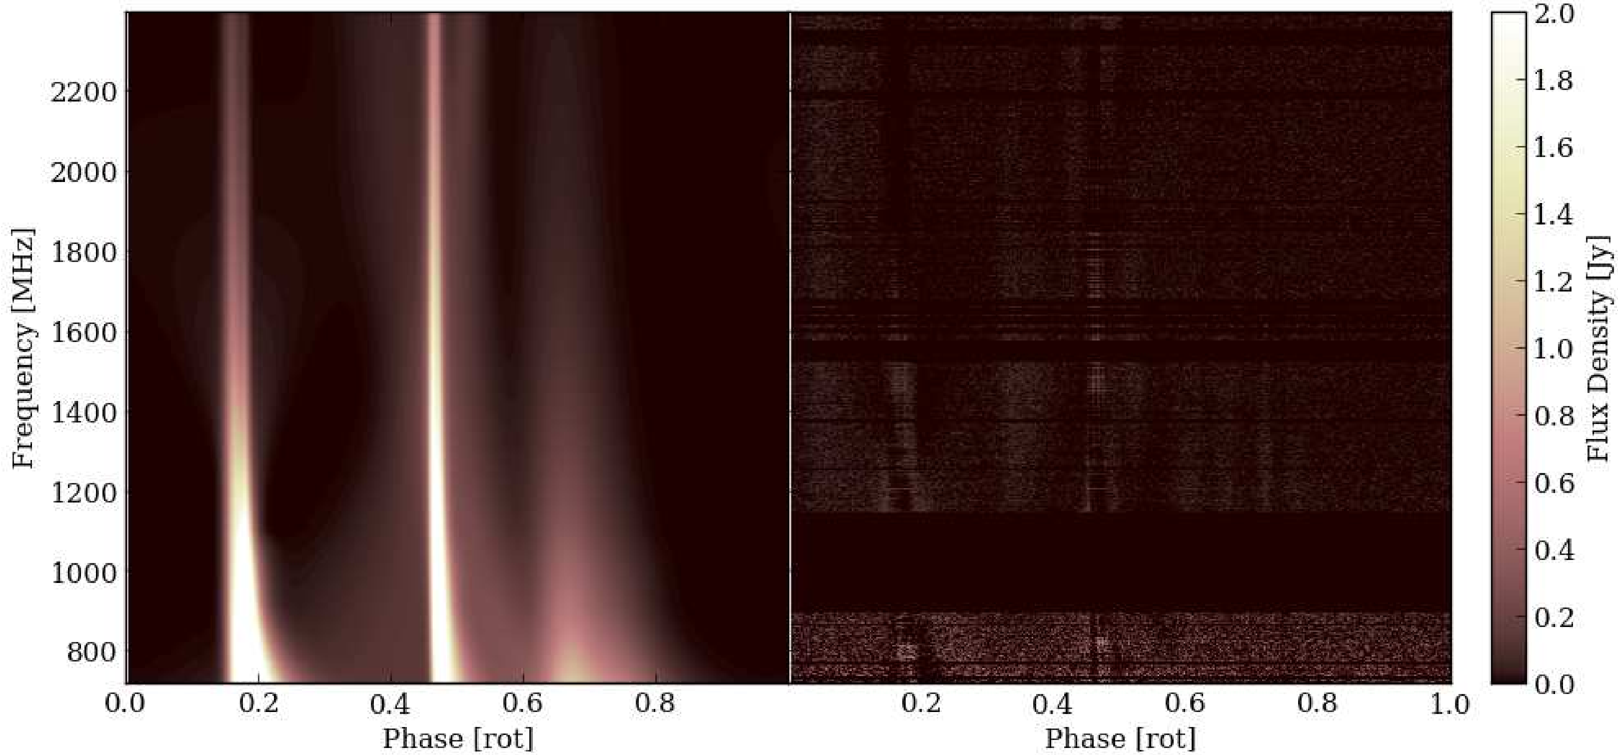

Fig. 3.3 – The constructed Gaussian model and residuals after subtracting the data in Figure 3.1. We used a ten-component model that captures both the finer structure seen at the higher frequencies and the scattering at the lower frequencies. Although the model and residuals show that the Gaussian modeling is not perfect, the model still proved sensible for timing and DM measurements. The phase-averaged spectral index of the model is consistent with that of Figure 3.1.

The scintillation bandwidth for M28A ($\sim$0.016 MHz at 1 GHz (Foster et al. 1991)) is much smaller than any of the observed channel bandwidths, so the data do not show obvious scintles in the folded profiles. However, to demonstrate the utility of fitting for the a_n parameters "for free", Figure 3.4 shows an example fit to fake data of moderate SNR generated by adding a fake "scintillation pattern" and frequency-independent noise to the M28A model in Table 3.2. A random phase and ΔDM was added to the data, and then it was run through our code, producing the fitted model and residuals shown in the figure. The fitted a_n values provide information about diffractive scintillation from the ISM, and they effectively act as weights for individual multi-channel TOAs that have been fitted for a DM and averaged together to obtain ϕ°_{ref}. Ideally, this advantage obviates the need to cull very low SNR TOAs of individual frequency channels. In principle, the a_n values could also be used to determine if there is residual RFI in the data, although we have not yet investigated how the presence of RFI will affect the fitting.

Using the algorithm described in §2.2.2, the Gaussian model was used as P_{nk} to fit for TOAs and DMs in the twenty-five observed epochs $\{D_{nk}\}$. The average per-epoch TOA uncertainty is $\sim$40 ns in the 1500 MHz data, and $\sim$90 ns in the 2000 MHz data. Figure 3.5 shows the measured DM variations for the M28A dataset, where an average DM of 119.88818 cm^{-3} pc was subtracted. We obtained DM precisions between several $\times 10^{-5}$ and several $\times 10^{-4}$ cm^{-3} pc. For the 1500 MHz data, the average DM precision of $\sim$7 $\times 10^{-5}$ cm^{-3} pc corresponds to about 140 ns $\approx 5 \times 10^{-5}$ rotations ≈ 0.02 bin of drift across the band, for 512 phase bins. It is interesting to compare this number to the amount of dispersive smearing in each channel from coherently dedispersing these data with the incorrect DM of 120 cm^{-3} pc; at 1500 MHz, $\delta t_{\rm DM} \approx 430$ ns. The first third of our measurements overlap with observations

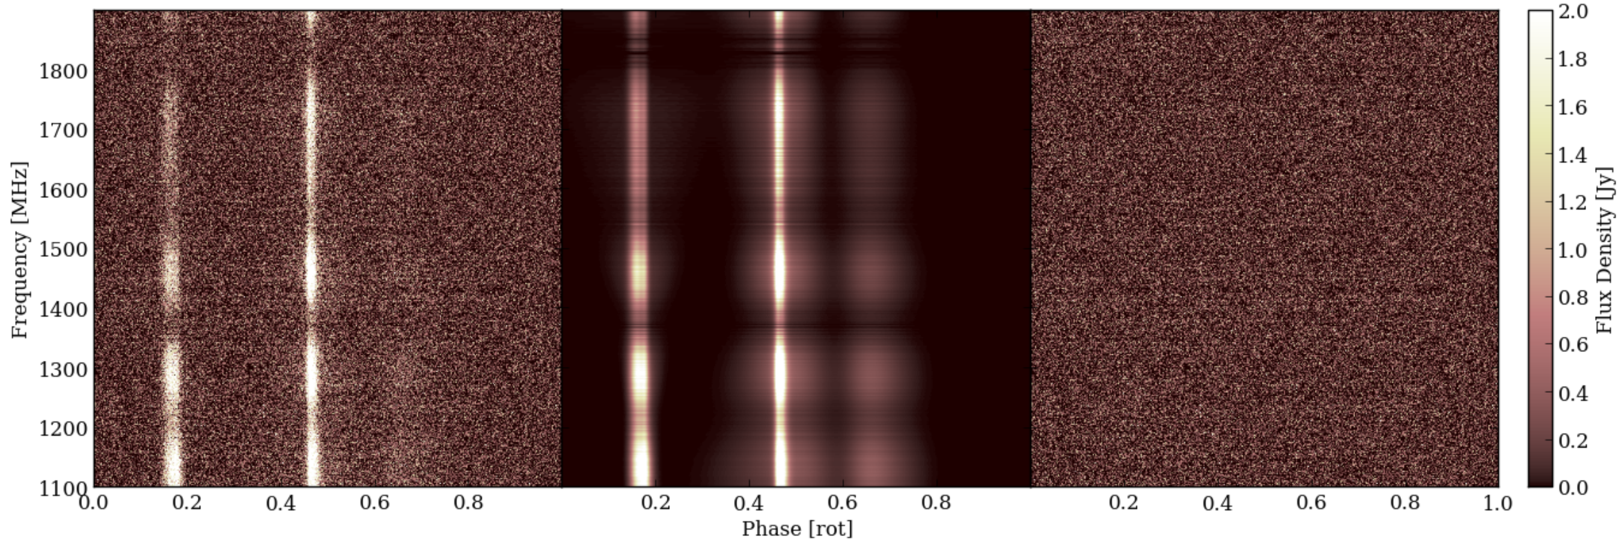

Fig. 3.4 – A simple demonstration of how the algorithm automatically manages diffractive scintillation. A random phase, DM, and "scintillation pattern" was added to the 1500 MHz portion of the model in Figure 3.3, along with frequency-independent noise (left panel), and then a fit was performed. The fitted model is shown in the middle panel, which encodes the scintillation in the a_n scale parameters, and the fit residuals are to the right. The residual statistics match the off-pulse noise and mean from the input data. A fake-data sample of Medium SNR from §3.2 has about the same SNR as these data.

of M28A presented in Keith et al. (2013); the overall trend in our DM measurements in these epochs is consistent with what is seen in their data.

3.1.2 Comparison of Methods

In what follows, we have compared our measurements (labeled "PP") with those obtained from the same data using a more traditional procedure (which we collectively label "TEMPO"). For the latter, multi-channel TOAs were obtained via standard techniques: each time-averaged epoch's band was divided into 64 channels, and each channel's profile was fit to a smoothed template profile that was obtained by averaging all the data from a given receiver. An FFTFIT-based algorithm was used for the pulse phase fitting [1]. Each epoch's DM was then determined by individually fitting a *fixed* timing model to the epoch's TOAs with `tempo`[2], *allowing only the dispersion measure* to vary[3]. In effect, this process fits removes a quadratic delay across the multi-channel TOAs. No consideration of profile evolution is taken into account besides the usage of three separate template profiles for the three bands. Therefore, if all of the TOAs were used in a timing model fit, the use of arbitrary phase-offsets (JUMPs) between the three sets of TOAs would be needed to align the template profiles.

Mitigation of Profile Evolution

The top panel of Figure 3.2 shows typical `tempo` multi-channel frequency residuals from not modeling the profile evolution. Note that the 820 MHz data is shown here to have the same channel bandwidth as the higher frequencies. Introducing phase-offsets to align small portions of the band (or from using numerous templates) is one

[1] Specifically, we used the Fourier phase gradient (PGS) algorithm in the `PSRCHIVE` program `pat`.

[2] `www.sourceforge.net/projects/tempo/`

[3] For clarity, at no time did we do a multi-band or multi-epoch fit for DM, although this is one area of current research.

approach to remove the frequency-dependent structure, but it adds a large number of otherwise meaningless parameters into the timing model (Demorest et al. 2013).

A somewhat less arbitrary approach is to characterize the trend with a simple function that can be included in the timing model. This latter multi-channel TOA strategy, combined with a `tempo` fit for the profile evolution and variable DM, is akin to the current timing methodology employed by the PTA collaboration called NANOGrav[4,5] (McLaughlin 2013; Zhu et al. 2015; Arzoumanian et al. 2015a). Ignoring profile evolution altogether and using frequency-averaged profiles may still be a sufficient practice for particular pulsars. However, this strategy will become untenable with the next generation of receivers. It seems more appropriate and simple to model directly the profile evolution based on the folded profiles, as we have done, and then simultaneously measure a TOA and a DM.

For the sake of comparison, we made multi-channel residuals *after* applying our algorithm to each of the same epochs in Figure 3.2, and have plotted them in the lower panel. These "residuals" were calculated by *independently* fitting each channel's profile in the *fitted* two-dimensional model with the corresponding profile in the data portrait using our own FFTFIT routine. The greatest improvement in modeling the profile evolution is seen in the 1500 MHz data, and we will show several consequences of this in the following sections. This improvement is sensible because the 1500 MHz data is our best "wideband" data in that it has the largest SNR, the largest fractional bandwidth, and hence the most profile evolution to be characterized. There is also continuity in the residuals with the 2000 MHz data (from a separate epoch). The

[4]The North American Nanohertz Observatory for Gravitational Waves: `www.nanograv.org`

[5]Here, we are refering to the use of a `tempo` functionality called "DMX", plus a polynomial function of log-frequency to account for profile evolution. The discrete DMs are measured in situ with other timing model parameters while using the overall WRMS residual as the discriminating quantity, which means the DMs and profile evolution parameter can absorb unmodeled, non-ISM effects like timing noise, or a gravitational wave signal.

2000 MHz data remains qualitatively the same because of its lower SNR and smaller fractional bandwidth (i.e. less observable profile evolution). The scatter of both sets of points is about the same as the corresponding average residual uncertainty.

On the other hand, the slight arch that remains in the 1500-MHz residuals and the added scatter into the 820-MHz residuals highlight the insufficiencies of a simple Gaussian modeling scheme for such a complex profile. The scatter in the 820 MHz points may be explained by the difficulty of characterizing its simpler profile with too many evolving Gaussian components, including scattering, although the fact that we did not apply any averaging to the model within each 3.125 MHz-wide channel may also play a role. More simply, an evolving Gaussian-component model does not describe the data well across all of the observed frequencies, but having more data in the 180 MHz-wide gap between the 1500 MHz and 820 MHz bands could help us find a better model. Note that these residuals are specific examples from the whole dataset, and that the goodness of the model's fit will vary from observation to observation.

Comparison of Dispersion Measures

The absolute DM is not a useful measure for comparison because its values depend on how the profile and its evolution are modeled, and the DM can even vary based on its inclusion in a timing model fit (not applicable here). Consequently, the average values of differently measured DMs will differ by a constant. Figure 3.6 shows that our mean-subtracted DMs are in agreement with those obtained from the above described methods. That is, the DMs measured in the 1500 MHz and 2000 MHz epochs are parallel to the solid line that represents equality, and so they track roughly the same amount of change in the DM.

The strong agreement in the DMs from the 2000 MHz data corroborates with our statement above about the similarity of the multi-channel residuals. The two sets of 2000 MHz ΔDMs agree within their errors and have a scatter of $\lesssim 2 \times 10^{-4}$ cm^{-3} pc $\approx$ 170 ns of drift across the band. In a similar vein, the observed scatter in the 1500 MHz data implies that our mitigation of the profile evolution mentioned above has significantly altered the measured DM trend. Here, the ΔDMs are scattered by $\lesssim 4 \times 10^{-4}$ cm^{-3} pc $\approx$ 800 ns of drift (with the largest deviation at three times that level). We address the measurement uncertainties in the next section.

The offsets seen in the figure between the DMs measured in each receiver band comes from the different modeling of profile evolution in each band. For example, the "TEMPO" DMs are measured with three different templates that are assumed to be constant as a function of frequency in their respective band. If this is a better assumption at higher frequencies, then the apparent average DM will be a function of frequency. Indeed, the pairs of observations that were separated by only $\sim$1 week and taken at different frequencies show an offset of $\sim 2 \times 10^{-3}$ cm^{-3} pc, which is much larger than any of the differences between DMs measured in the same band, on the same time scale.

Similarly, having tried a vast number of fitted Gaussian models for M28A, we found that switching between different families of models produced the same large offsets (up to a few $\times 10^{-3}$ cm^{-3} pc) between the DMs measured in different bands. Given the reasonable assumptions we made about our model, we believe that the large frequency-dependent offsets seen in the DMs measured by using other fitted models (and the "TEMPO" templates) is explained by a misrepresentation of M28A's profile evolution and not, for instance, a frequency-dependent DM. In fact, we used the assumption that no such offset exists in temporally proximate data from different frequencies as a qualitative model-selection criterion.

Ultimately, dispersion measures will be a function of frequency since the multipath propagation of different frequencies will sample slighty different total free-electron column densities (Cordes et al. 2015). However, there is ambiguity between a frequency-dependent DM and profile evolution; as noted by Ahuja et al. (2005, 2007), Hassall et al. (2012), and others, an apparent frequency-dependent DM can be explained by unmodeled profile evolution. For example, again consider Figure 3.2; the DM measured using different sets of frequencies would vary because the phase-offset between arbitrarily chosen pairs of frequencies is not constant.

Potentially, in a bright, highly-scattered, high-DM pulsar like M28A, a frequency-dependent DM could be detected. A rough estimate of the level of $\delta\mathrm{DM}(\nu)$ that could be expected in the data can be garnered from estimates of M28A's scattering measure and distance as reported in Foster et al. (1991) in combination with the prediction for the form of a frequency-dependent DM in §4.4 of Cordes & Shannon (2010). However, the prediction is nearly proportional to the unknown distance to the scattering material. Furthermore, a constant offset between DMs determined in different frequency bands is highly covariant with profile shape evolution, as described above.

It may become feasible to disentangle a frequency-dependent DM from profile evolution when truly broadband (eg. fractional bandwidth $\gtrsim$1), long-term observations become readily available, since having both temporal and frequency DM variations can break the degeneracy. We will save the detailed question of a frequency-dependent DM for future investigation, as it is important for those who will correct high frequency data with DMs measured at low frequencies. Although we offer no solution to the problem of disentangling profile evolution and dispersion measure, we can give greater credibility to measurements of dispersion measure changes, which are the more important quantities for timing experiments, and perhaps more interesting for studies of the ISM.

Comparison of Measurement Uncertainties

Figure 3.7 shows a comparison of the uncertainties on the TOAs (left panel) and DMs (right panel). The TOA uncertainties shown here are from frequency-averaged TOAs (one TOA per epoch, per band) that were obtained in a similar fashion as the multi-channel TOAs. The DMs measured from the multi-channel TOAs were used to align the profiles before frequency-averaging them; this is one traditional way of accounting for significant DM changes, which would otherwise smear the average profile and systematically inflate the TOA error by an amount related to ΔDM. Each of these TOAs will reference some specific frequency and will be covariant with the DM. In order to make a fair comparison, we have plotted the transformed "PP" TOA uncertainties to reference these frequencies; the zero-covariance uncertainties are smaller by $\lesssim 20\%$.

However, we did not use a weighted frequency-average, neglecting any SNR variation across the band that might originate from ISM effects or profile evolution; this is a second effect that can lessen the timing precision in the standard protocol. Weighting the multi-channel TOAs (or the pulse profiles) to obtain a single frequency-averaged TOA reduced the "TEMPO" TOA uncertainties in Figure 3.7 by factors between one and three, bringing all of them to within a factor of two of the "PP" uncertainties.

Finally, systematic trends from profile evolution (see Figure 3.2) will enlarge the uncertainties. The effect on the timing and DM precision from marginalizing M28A's profile evolution is unambiguous in the 1500 MHz data; the TOA uncertainties ob-

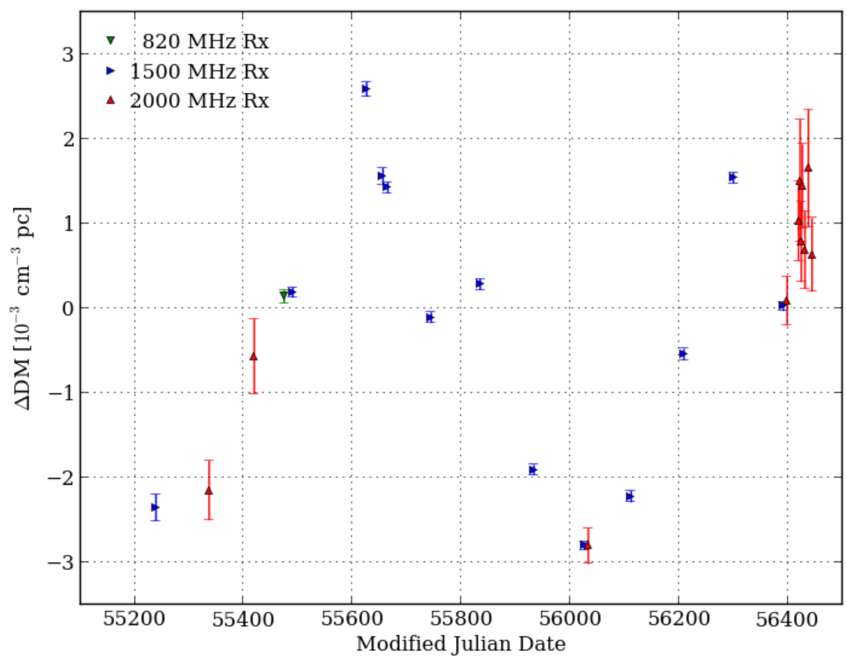

Fig. 3.5 – The twenty-five DM measurements from Table 3.1. The calendar range of the data spans Feb 11, 2010 to May 31, 2013. Keith et al. (2013) reports a similar $\sim$5 $\times 10^{-3}$ cm^{-3} pc increase in the first third of our data.

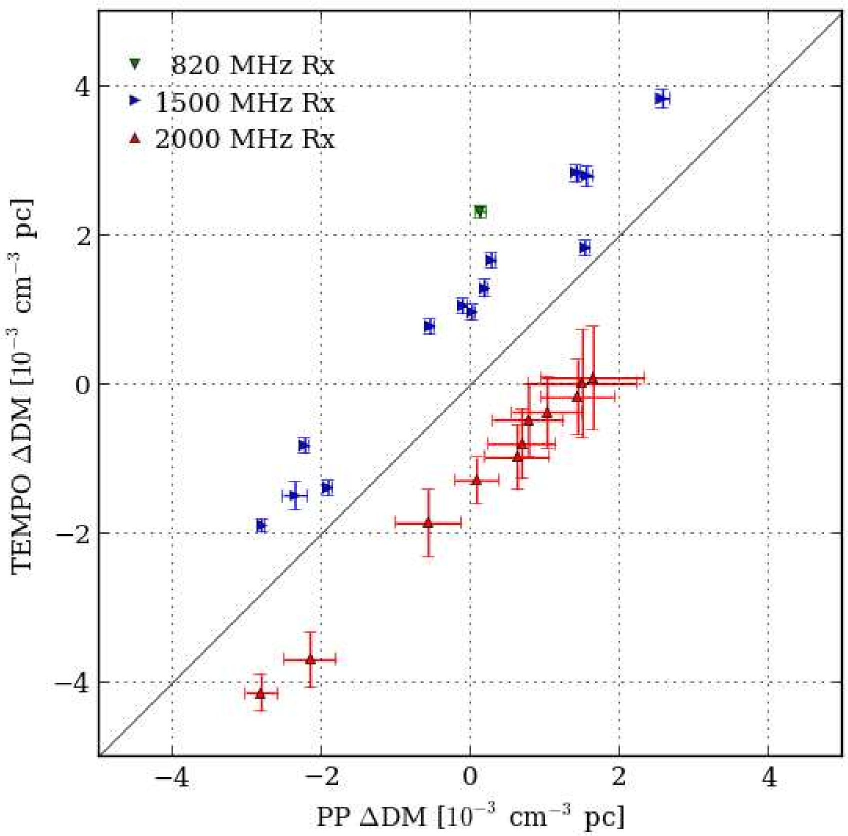

Fig. 3.6 – A comparison of mean-subtracted DM trends as measured in the M28A data by our technique and a more usual approach. "PP" represents our measurements. The solid line traces equality. The error bars for the "TEMPO" DMs are given by the least-squares `tempo` fit, whereas the calculation of the error on our DM measurements is provided in §2.3. The $\sim 2 \times 10^{-3}$ cm^{-3} pc offset arises from the difference in how profile evolution is modeled. See text for further discussion.

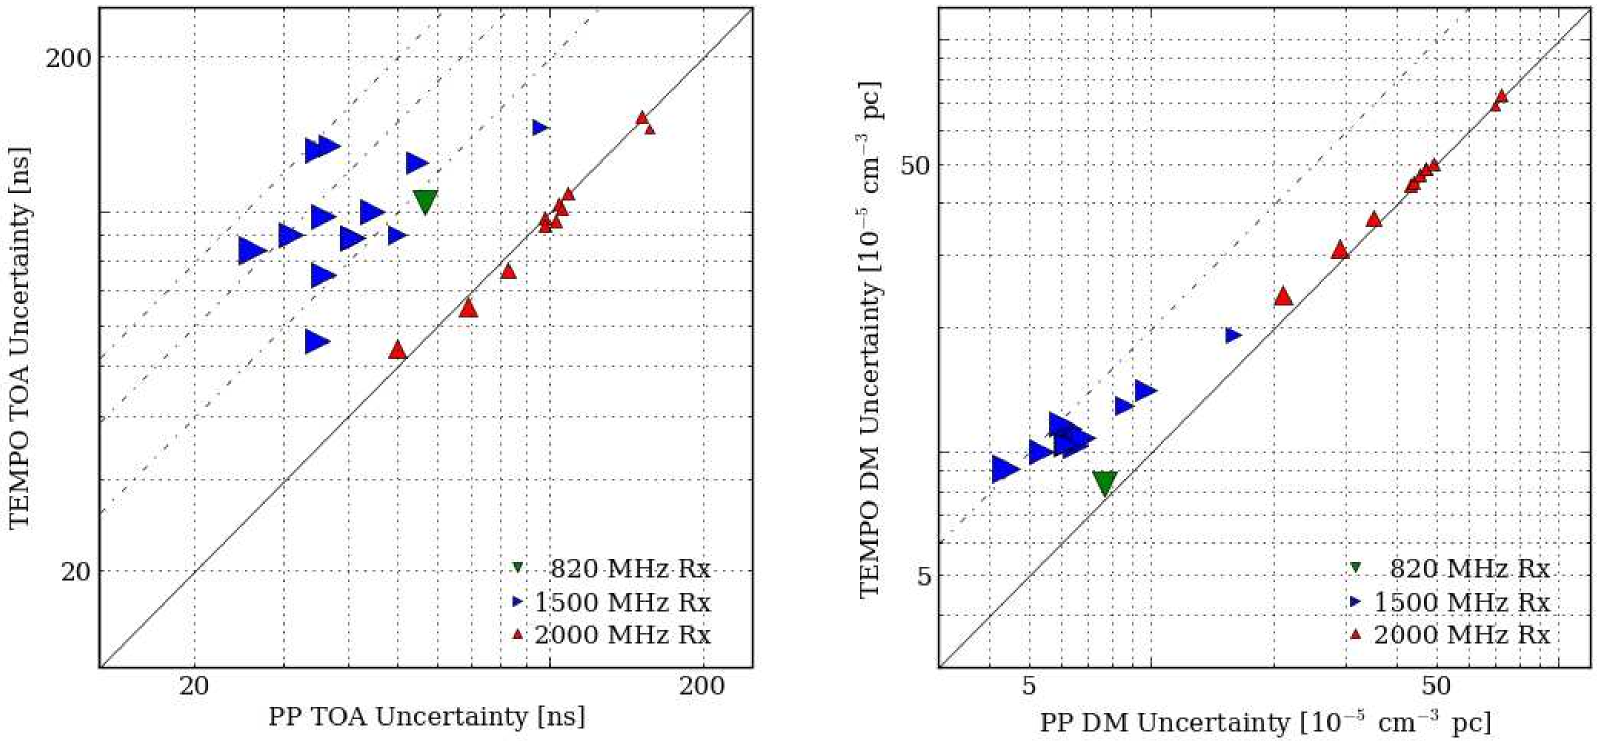

Fig. 3.7 – A log-scale comparison of the TOA and DM uncertainties (left and right, respectively) from the M28A data. The dash-dot lines indicate differences by factors of two (both panels), three, and four, and the area of each triangle is proportional to the data's SNR. The largest improvements are seen in the 1500 MHz data, where profile evolution has been most mitigated. The "PP" TOA uncertainties have been transformed to the same set of reference frequencies. Two points in the left plot have the same values, so it appears as though there are only twelve 1500 MHz epochs.

tained by the new algorithm using a simple Gaussian model are smaller by up to a factor of four, with an average of about five-halves, and the DM uncertainties are smaller by up to a factor of two, with an average of about three-halves. There is no improvement in the 2000-MHz TOA uncertainties, where there is no significant profile evolution, and the improvement of the 2000 MHz DM uncertainties is marginal, but may be a function of SNR.

The dependency of the improvement on the signal-to-noise ratio is expected because any correctable profile evolution becomes evident with increasing SNR. This is particularly evident for the DM uncertainties, whereas the 1500 MHz "TEMPO" TOA uncertainties are dominated by systematic error from averaging over the profile evolution. A comparison of the uncertainties as a function of SNR shows that they are all roughly proportional to the SNR, except for the 1500 MHz "TEMPO" TOA uncertainties.

The ad hoc methods to mitigate effects arising from dispersion measure changes, frequency-dependent SNRs, and profile evolution in wideband data, are all naturally accounted for by using the new algorithm, which we have seen to yield superior, or at least as good, measurement precisions.

3.2 Monte Carlo Analyses

3.2.1 Description

We completed a variety of Monte Carlo analyses to explore the accuracy to which the algorithm can determine parameter estimates, errors, and covariances in a number of regimes, which included varying the data resolution, the signal-to-noise ratio, scintillation patterns, and level of δDM. Here, we show results from generating fake pulsar data by adding random, frequency-independent noise to the total-intensity model for M28A given in Table 3.2. The data emulate those from typical pulsar timing observations with GUPPI; we set the center frequency to 1500 MHz and the bandwidth to 800 MHz.

We explored the performance of the algorithm in a variety of data resolutions, changing the number of phase bins in the profile from 128 to 2048, and the number of frequency channels from 8 to 512, both in powers of two. When using an insufficient number of phase bins to resolve the profile, some harmonic power gets aliased into the estimate of the profile's noise level, which in turn suppresses the estimate of the parameter errors. Since we expect these issues to be avoided in practice, and because our results seemed independent of the number of profile bins once the profile is resolved, we restrict the presentation of the Monte Carlo trials to those with profiles having $n_{bin} = 2048$.

For each sample in a given Monte Carlo trial, a random infinite-frequency phase was drawn uniformly from the interval [-0.5, 0.5) and injected into the model. The injected DM value was the nominal ephemeris value plus a perturbation drawn uniformly from the $\log_{10}$ interval of approximately [-5.0, -1.5], with equal probability given to the sign of the perturbation. We chose this interval because it equally samples different scales of perturbations with a maximum that roughly corresponds to $\frac{\Delta\mathrm{DM}}{P_s} \approx$ 100 [10^{-4} cm^{-3} pc ms^{-1}], and because we do not expect in most cases that a DM will be determined to better than $\sim10^{-5}$ cm^{-3} pc. For simplicity, at a given SNR the RMS noise level remained constant as a function of frequency across the band, but in some of the tests we added random amplitude patterns (a_n) to mimic the effects of scintillation (not presented here, but see Figure 3.4 for an example).

3.2.2 Results

Figures 3.8, 3.9, and 3.10, show results from the Monte Carlo trials in three SNR regimes for the seven values of n_{chan}. We used `PSRCHIVE` to measure the noise level in the data, and the three SNR levels in our trials presented here were set to be near the `PSRCHIVE` values of 20.0 ("Low"; yellow, dash-dot), 100.0 ("Medium"; purple, dashed), and 1000.0 ("High"; gray, solid)[6].

Figure 3.8 shows two statistics returned from the trials for the phase estimates (left column, squares) and DM estimates (right column, diamonds) as a function of n_{chan}. The top row shows the mean of the distribution of the values

$$\frac{\text{estimated value} - \text{injected value}}{\text{calculated error}}, \tag{3.1}$$

and the bottom row shows the standard deviation of this distribution. If there are no systematic differences and if the errors are calculated accurately, then this distribution should be $\sim$ Normal$(0, 1)$. There is no obvious evidence of bias as a function of n_{chan} or SNR, meaning the injected values are accurately recovered, within the error. Even though all of the recovered normalized distributions were very well approximated by a Normal distribution (down to very small SNRs), one can see that the errors are underestimated when the channel-SNR becomes sufficiently low. However, even in the Low SNR case with the largest number of channels, the errors are off by no more than 20%.

Figure 3.9 shows how the absolute errors change with n_{chan}. We have separated the trials for clarity; each point represents the median of the error distribution, contained within the 95% highest-density region. The uncertainty scales linearly with the SNR for both parameters. There is no obvious dependence on the average TOA error with

[6] The SNRs of the M28A data presented in the previous section varied between $\sim$100 and $\sim$700.

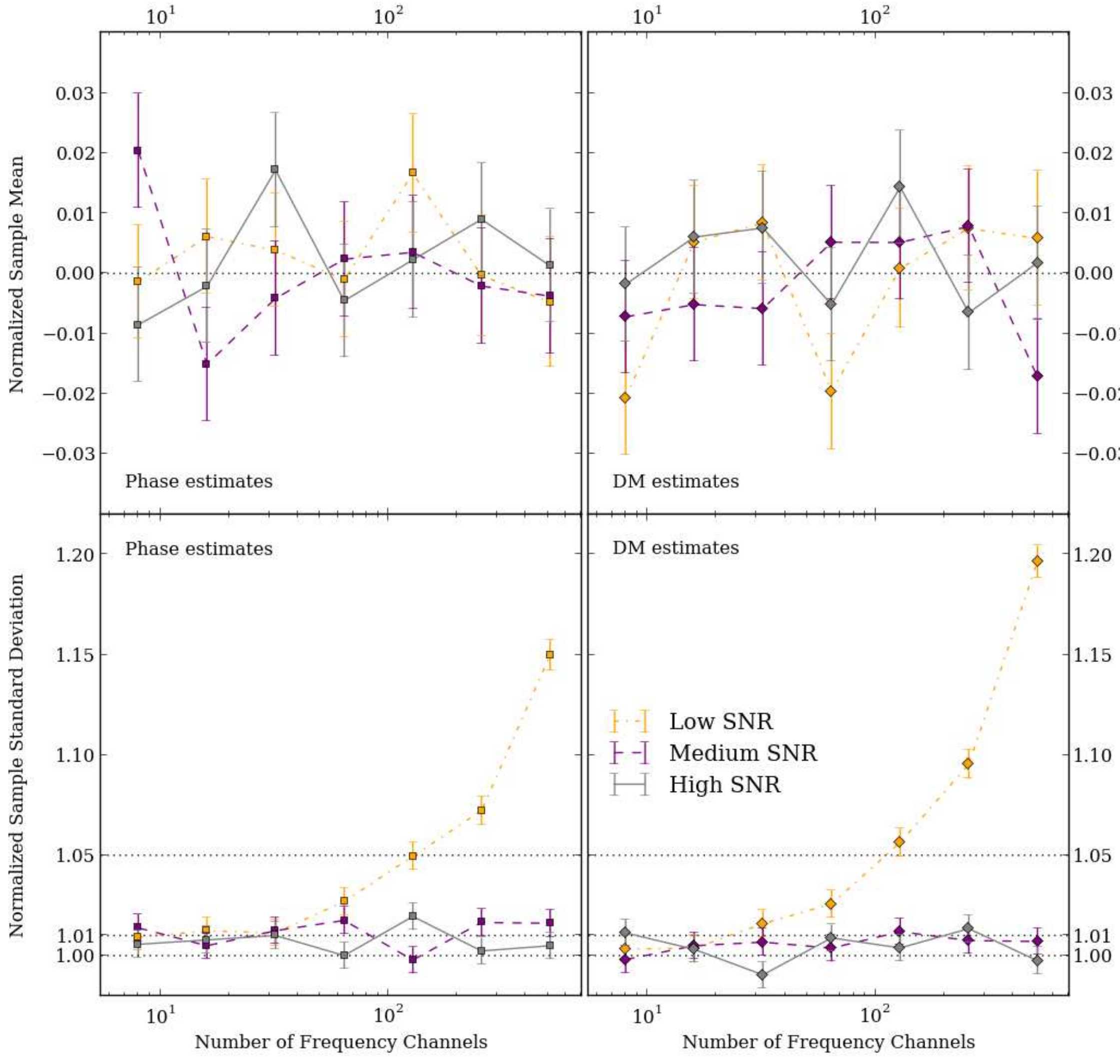

Fig. 3.8 – Monte Carlo results for examining bias and error calculation as a function of n_{chan} and SNR. Each trial consisted of 11,400 samples. The squares (left column) show results in each of three SNR regimes for the phase estimates, and the diamonds (right column) show the same for the DM estimates. The two statistics shown are the normalized sample mean (top row) and standard deviation (bottom row). The error bars are each one standard error of their respective statistic. The dotted lines in the bottom row correspond to a 0%, 1%, and 5% underestimation of the errors. Additional Monte Carlo trials were performed for a wider range of SNRs, which fill the gap between the Medium and Low SNRs shown here, as well as perform even more poorly than the Low SNR trial.

n_{chan}, but there is some increase in the DM error by a few percent as the number of channels becomes small, independent of the SNR. This is expected after considering Equation 2.22b because the effective frequency range of the data (the difference of the center frequencies in the highest and lowest channels) decreases with the number of channels as n_{chan}^{-1}; that is, the "lever arm" for the DM measurement lessens, giving a greater measurement uncertainty.

The non-Gaussianity of the error distributions in Figure 3.9 is somewhat noticeable in the Low SNR regime towards higher n_{chan}, but it is manifested in the trend seen in the lower half of Figure 3.8. The skewness towards small uncertainties becomes very obvious at lower SNRs (not shown here). The underestimation of TOA errors at low SNRs (particularly for profiles with large duty-cycles) has been documented before (Hotan et al. 2005).

Lastly, there is a dependency of each error distribution's variance on n_{chan}. Beyond some low channel-SNR, the variance appears constant, which has been verified for the Middle SNR case. The variance of the error distribution is not a particularly interesting quantity, so we will refrain from additional discussion, making only a note that it seems to affect the value of the errors at the level of a few percent. After replicating this series of Monte Carlo trials with a similar Gaussian model that has no profile evolution (i.e. the same components with constant positions, widths, and amplitudes), we find almost identical results, except that the absolute errors uniformly decreased by $\sim$3%. From this we conclude that the effects from marginalizing over profile evolution in each channel for this model were minimal.

It is important to remember that all phase and error estimates shown here are referenced to ν_{zero}; if a different reference frequency were used, the results in Figures 3.8 and 3.9 would look different because of non-zero covariance between ϕ°_{ref}

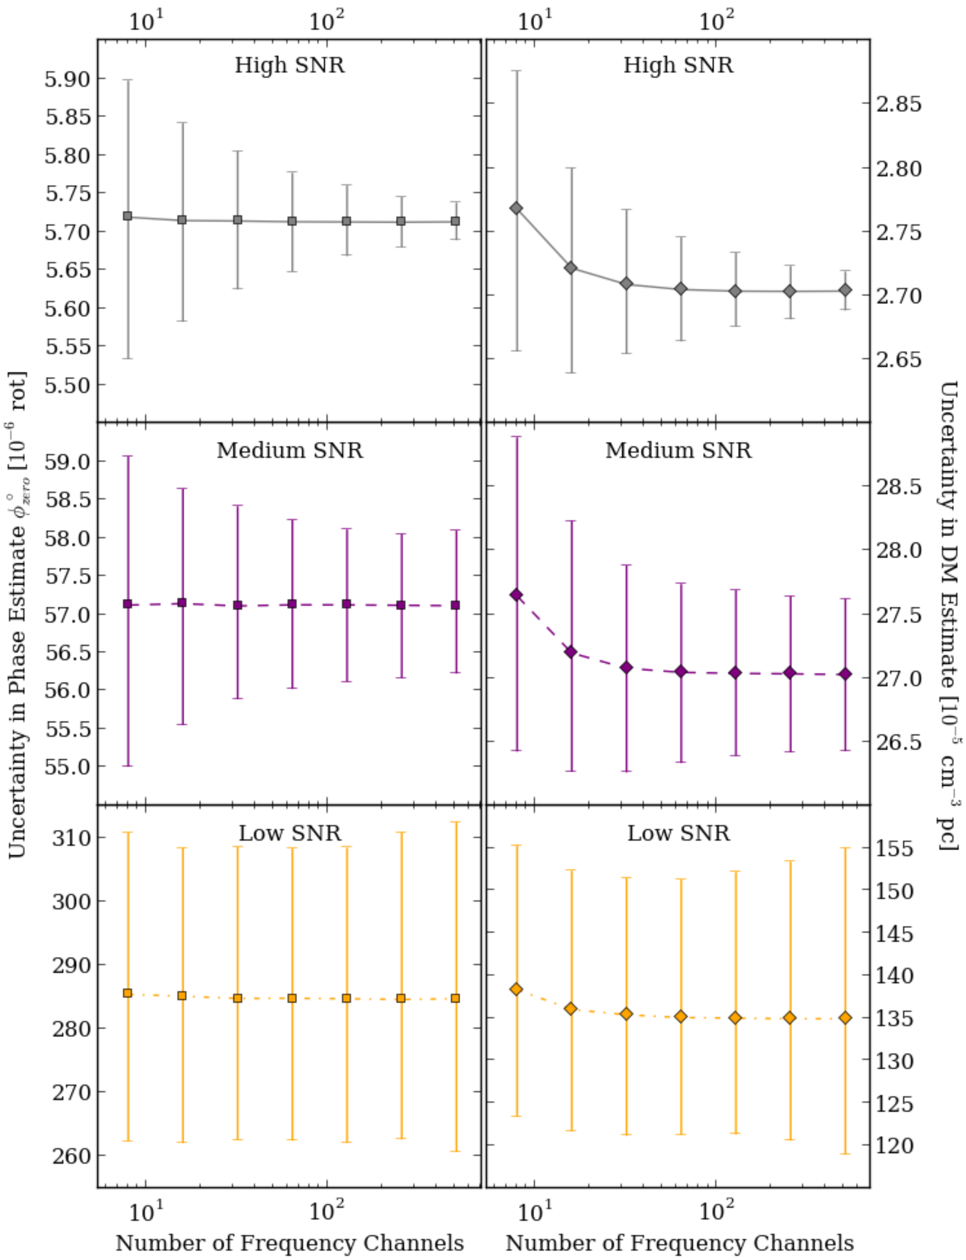

Fig. 3.9 – The error distributions' dependency on n_{chan} is plotted for the three SNR regimes. The plotted points show the distributions' median values and 95% highest-density regions. There is a slight skew in the error distribution for the Low SNR regime that becomes conspicuous at even lower SNRs. Additional Very Low SNR Monte Carlo trials not shown here have severely skewed error distributions.

and DM. To verify if the calculated ν_{zero} is, in fact, the zero-covariance dedispersion reference frequency, we show the sample correlation coefficient (the sample covariance normalized by the sample standard deviations) in Figure 3.10 for the same three sets of Monte Carlo trials. It is obvious that there is non-zero covariance in the Low SNR regime.

One underlying issue that can explain this feature is that our analytic formulation for ν_{zero} given in Equation 2.24 will not be precise for all data resolutions of arbitrary SNR. This is verified for the low SNR case in Figure 3.11, which shows the discrepancy between our calculated covariances and the covariances measured in additional Monte Carlo samples. The Monte Carlo trials are the same as before, but are now fixed with $n_{chan} = 512$, while varying ν_{ref}. The vertical dotted line shows the calculated ν_{zero} for this SNR, which is significantly offset (~30 MHz) from the interpolated zero-crossing of the sample correlation coefficient curve. That the slopes of the functions in Figure 3.11 are steepest near the zero-crossing implies that the determination of the parameter uncertainties is sensitive to the determination of ν_{zero}.

3.3 Conclusions

We have presented a novel but simple routine for the simultaneous measurement of TOAs and dispersion measures in folded pulsar data by using a frequency-dependent model of the pulse profile. This algorithm is a straightforward extension of FFTFIT from Taylor (1992), but it has some advantages over more standard techniques; these include being able to directly characterize ISM effects (such as dispersion measure changes, scintillation, and scattering) and profile evolution, while removing the need for additional, undesired parameters in the timing model (i.e. JUMPs). Any arbitrary model can be used, but the choice of model will affect the measured values.

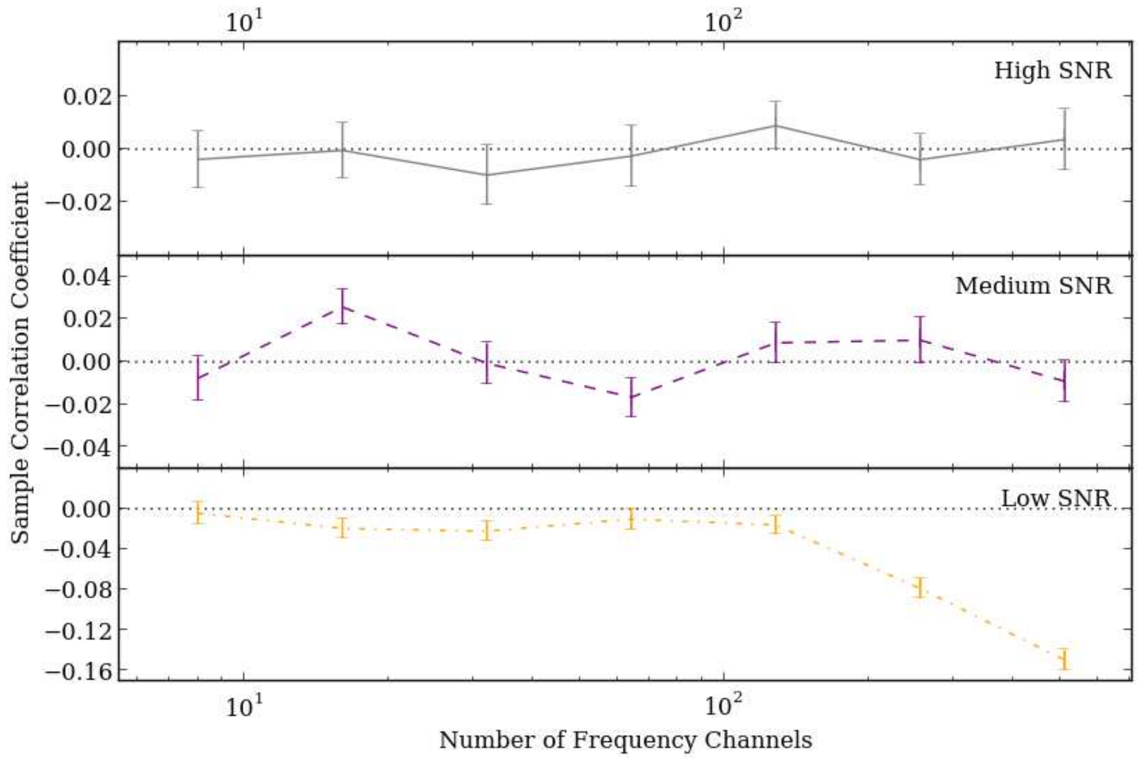

Fig. 3.10 – The samples' normalized covariances are plotted to verify that ν_{zero} is the zero-covariance reference frequency. The point estimates for the sample correlation coefficients and their errors were determined by a resampling analysis of each full Monte Carlo trial. The trend prevalent in the Low SNR regime is due to an inaccurate determination of ν_{zero} (see text and Figure 3.11).

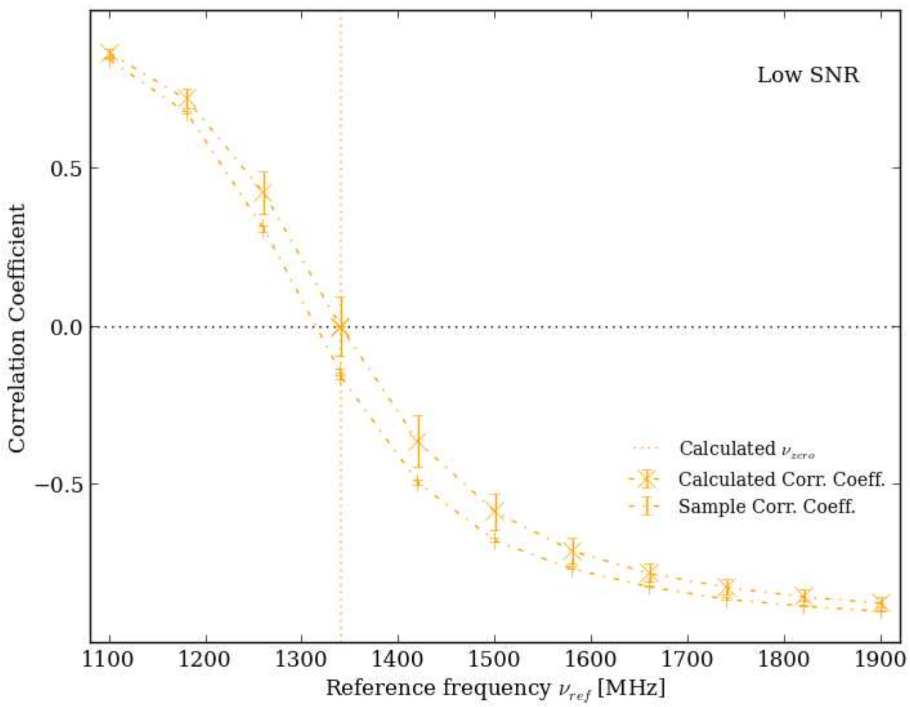

Fig. 3.11 – The Low SNR correlation coefficient for Monte Carlo trials as a function of ν_{ref} for $n_{chan} = 512$. The vertical dotted line shows the calculated value for ν_{zero} (the average value is plotted, but there is effectively zero dispersion). The interpolated value for the "true" ν_{zero} in the sample differs by ∼30 MHz. Note that the sample correlation coefficient at ν_{zero} here agrees with that from Figure 3.10 for $n_{chan} = 512$. The equivalent curves for the higher SNR trials overlap almost exactly. The point estimates of the sample correlation coefficients and their errors were determined in the same way as Figure 3.10.

The demonstration of a simple Gaussian modeling scheme to make these measurements in a 3-year, wideband dataset of the millisecond pulsar M28A shows that we are able to obtain reliable measurements of the dispersion measure, as well as improved TOA and DM precisions by up to a factor of four for the former and two for the latter. The biggest improvements in the parameter precisions and in mitigating profile evolution were seen in the high signal-to-noise 1500 MHz data, which was our best "wideband data" for demonstration purposes in the sense that it has the largest fractional bandwidth (and therefore the most obvious effects from interstellar dispersion, scattering, and profile evolution). Similar improvements in other pulsars will depend on the mitigation of intrinsic and extrinsic profile evolution. It became clear in our comparisons that there is a necessity for quantitative model selection based on more robust two-dimensional portrait modeling, which potentially can lead to the detection of a frequency-dependent DM, or other interesting signals.

The results from our Monte Carlo analyses has led us to the conclusion that a large number of frequency channels is appropriate for applying this technique. A larger number of channels will provide the highest precision DM measurements and avoids averaging over profile evolution. The proper incorporation of discrete DM measurements with their own heteroskedastic errors (besides the TOAs') into the determination of a timing model (eg. by using `tempo`) is not trivial, but a Bayesian approach has been investigated in Lentati et al. (2013). Relatedly, measuring DMs from non-simultaneous, but temporally proximate multi-frequency data can be an intermediate improvement until larger bandwidths become readily available. This is another avenue of future development, although it comes with the drawback of having correlated TOAs.

One important caveat in these Monte Carlo tests is that the model fitted to the simulated data was the true model from which the data were generated. In practice, a Gaussian-component model fitted to real data will not match perfectly (leaving behind non-Gaussian residuals in all but the simplest or low SNR cases) and there will be a much stronger dependence of the measured DM and TOA on the number of channels and the amount of profile evolution. This is one area requiring further testing, but it also suggests to err on the side of more frequency channels.

Although general, the algorithm will be most useful when applied to MSPs because of their sensitivities to small dispersion measure changes, as highlighted by Equation 2.3, and because of the need to correct for their profile evolution in wide-band data to obtain the highest possible timing precisions. For these reasons, we believe this algorithm will provide a natural TOA and DM measurement procedure for campaigns of MSP monitoring, like that of NANOGrav or other PTA experiments.

Chapter 4

Application of Wideband Timing to NANOGrav MSPs

Note: The results from this chapter will be published as a complimentary analysis of the data recently published by NANOGrav: "NANOGrav Nine-Year Data Set: I. Arrival Time Measurements and Analysis of 37 Millisecond Pulsars", Arzoumanian, Z. et al. (2015); arXiv:1505.07540.

Abstract

In this chapter, we apply the methods from Chapters 2 and 3 (i.e., Pennucci et al. 2014) to a subset of 37 millisecond pulsars (MSPs) currently observed by the North American Nanohertz Observatory for Gravitational Waves (NANOGrav). The future of pulsar timing array experiments will include very broadband receiver systems, the data from which will require novel analysis techniques. The current capabilities of NANOGrav lie squarely between dated and mature technologies. NANOGrav makes wideband (hundreds of MHz) observations in a number of frequency bands per epoch per pulsar. While suboptimal for the ideal wideband timing approach, we can still use our methods in tandem to the current, standard NANOGrav efforts and see how they compare. Recently, NANOGrav submitted for publication a paper containing the TOA and timing model data products obtained from the underlying profile data that we analyze here (Arzoumanian et al. 2015a). Our analyses of this 9-year data set are entirely parallel to those presented in Arzoumanian et al. (2015a), and so throughout this chapter we refer to the current NANOGrav techniques and draw comparisons. We address profile evolution by building Gaussian component model portraits for all pulsars, and then make timing measurements, and perform a full Bayesian analysis of the noise. We highlight both pitfalls and successes in the current iteration of our wideband method: while Gaussian modeling is far from perfect, we obtain consistent timing results in most cases, and even mitigate specific chromatic red noise residuals. The results from this project will later be published as a NANOGrav paper, and the present results will be analyzed in parallel with the Arzoumanian et al. (2015a) data to place the most stringent limits on the stochastic gravitational wave background in the coming months.

4.1 Introduction

One of the primary motivations behind the work of Chapters 2 and 3 was to improve the timing of millisecond pulsars (MSPs) for the sake of gravitational wave detection. By "improve the timing", we mean to ultimately improve the precisions of the pulse times-of-arrival (TOAs) and the variable dispersion measure (DM), which in turn will admit better estimates of the timing model parameters. These improvements should, in principle, increase the sensitivity to gravitational waves, if not just marginally. As discussed earlier in Chapter 2, these improvements are accomplished by mitigating the effects of profile evolution from the interstellar medium (ISM) and the pulsar itself.

A second, more pragmatic motivation for our developments is to simplify the data sets; the current methods in NANOGrav yield a single TOA per frequency channel per epoch, which constitutes an undesirable amount of TOAs (∼50–100 per pulsar per epoch). Profile evolution is modeled as a time-constant set of *ad hoc* parameters and the DM as a piece-wise constant function, both of which are included as part of the timing model. Our methods measure a single TOA and a single DM per receiver per epoch, where in the future we envision only a single receiver will be necessary to sufficiently constrain the DM. Our pulse profile evolution model is fixed in a two dimensional phase-frequency "pulse portrait" in a manner analogous to how one dimensional standard template profiles are fixed for the generation of channelized TOAs. Although a model for DM(t) will still be necessary as part of the timing model, our DM measurements provide, ideally, anchor points for the model. In the most banal scenario where either method — using channelized TOAs, or single wideband TOAs and DMs — yields precisely the same timing results, the reduced data volume of the wideband data significantly decreases the computational complexity of gravitational wave searches[1], which scale approximately with the number of TOAs squared (J. Ellis, private communication; Arzoumanian et al. 2014, and see the forthcoming NANOGrav detection paper based on the 9-year data set.).

Of course, a reduced data volume in this case could be arrived at by appropriately averaging the channelized TOA data, but a third objective of the wideband approach is to provide astrophysically interesting parameters from segregating intrinsic and extrinsic profile evolution (as opposed to, e.g., the "FD parameters" described in §4.1.1), be they measurements of interstellar scattering manifested as pulse broadening, or the spectral dependencies of components. We will briefly discuss this topic in §4.2.1. However, we note that some results based on this motivation have been or will be published (Bilous et al. 2015; Pennucci et al. 2015, or see Chapter 6).

4.1.1 NANOGrav: The Nine-Year Data Set

A brief introduction to the North American Nanohertz Observatory for Gravitational Waves (NANOGrav; GWs) as a pulsar timing array (PTA) collaboration is given in §1.4.1. The data set we concern ourselves with here is a largely homogeneous set of observations spanning from approximately 2005 to 2014, for the longest observed pulsars. The data set builds on the 5-year data set analyzed in Demorest et al. (2013), which presented the first NANOGrav limit on the stochastic background. The TOAs and fitted timing models determined from the folded light curves (the "profiles") of this extended data set, which we will refer to as "the 9-year data set", was recently submitted for publication (Arzoumanian et al. 2015a, hereafter A15). Gravitational wave analyses of these data by NANOGrav will follow in the coming months and will follow up on the limits presented in Demorest et al. (2013), Arzoumanian

[1]Since GWs affect all radio frequencies equally, the GW searches do not explicitly care about the frequency of the TOAs; all TOAs get referenced to infinite frequency.

et al. (2014), and Arzoumanian et al. (2015b). We direct the reader to Demorest et al. (2013) and A15 for a detailed presentation of these data; here we will present a brief overview of the observing program and the data set.

Pulsars & Observing Program

NANOGrav carries out its observing program with the 100-m Robert C. Byrd Green Bank Telescope (GBT) of the National Radio Astronomy Observatory and the 305-m William E. Gordon Telescope of the Arecibo Observatory (AO); any NANOGrav pulsars within the declination range ($0^\circ < \delta < 39^\circ$) are observed at Arecibo, with the rest covered by the fully-steerable GBT. At the time of writing, NANOGrav observes 49 MSPs: 28 at Arecibo and 23 at Green Bank; MSPs J1713+0747[2] and B1937+21[3] are observed at both. Because GW limits from PTAs are dominated by the best pulsars (see, for example, Demorest et al. 2013; Arzoumanian et al. 2014), it is sensible to observe them with both telescopes; they also provide precise consistency checks between the various observing setups at the two telescopes (they are both observed in the three higher frequency bands). 37 of these 49 MSPs are presented in the 9-year data set, which is 20 more than in the 5-year data set; the increase in MSPs is primarily a reflection of numerous successful surveys directed at finding PTA-quality MSPs through either blind or directed searches (Ransom et al. 2011; Ray et al. 2012; Boyles et al. 2013; Lynch et al. 2013; Stovall et al. 2014; Lazarus et al. 2015). NANOGrav expects to continue to add MSPs into the array at a rate $\sim$4 yr^{-1}. According to Siemens et al. (2013), adding MSPs to a PTA is the best way to increase sensitivity to a stochastic background after the

[2] J1713+0747 is the best long-term-timed MSP in the northern hemisphere and was the subject of a recent NANOGrav timing group paper that presented 21-year timing results (Zhu et al. 2015).

[3] B1937+21 was the first millisecond pulsar discovered (Backer et al. 1982) and is a landmark in pulsar astrophysics; it remains the fastest field MSP known (rotating $\sim$642 times a second), beaten only by the MSP J1748−2446ad in the globular cluster Ter5a (rotating $\sim$716 times a second).

lowest frequency GWs begin to dominate the residuals. It should be noted that a significant quantity of archival data exists for a number of MSPs that will serve to extend NANOGrav baselines by 5–15 years when integrated into the data set.

Basic parameters and observing spans for the 37 MSPs in the 9-year data set are given in Table 4.1. The current observing program has a cadence of approximately 1 observation per pulsar per frequency band per 3–4 weeks[4]. Each observation is approximately 20 min in duration, during which time all four polarization products are recorded for full Stokes recovery. In addition to each source scan, a calibration scan is taken, which consists of a short on-source observation while a noise diode is pulsed at 25 MHz. The same noise diode is pulsed on and off of a flux calibrator source (a quasar) approximately once a month. These sets of calibration scans allow for full flux and polarization calibration of the data. The latter is more important for high precision timing, as improper weighting of the polarization channels can induce spurious profile shape changes, leading to bias in the TOAs (van Straten 2013). Note that our work, however, only considers the calibrated total intensity measurements. In total, NANOGrav averages on the order of ∼1-2 hr day^{-1} on each telescope[5].

For each pulsar, measurements are made in two disparate frequency bands in order to precisely measure the DM and remove the dispersive time delay relative to an infinite frequency signal (see §2.1),

$$\delta t_{\mathrm{DM}}(\nu) = K \times \mathrm{DM} \times \nu^{-2}, \tag{4.1}$$

for frequency ν and where $K^{-1} \equiv 2.41 \times 10^{-4}$ MHz^{-2} cm^{-3} pc sec^{-1} (Lorimer &

[4]This program is augmented by ∼weekly observations of some of the best pulsars at either telescope in order to increase sensitivity to continuous wave sources.

[5]With limited trained personnel, this currently amounts to a bare minimum of ∼50 hr of time per observer per year, including the author.

Table 4.1. NANOGrav Pulsars

Source	P [ms]	$\dot{P}$ $[10^{-20}]$	DM [pc cm^{-3}]	P_b [d]	Span [yr]
J0023+0923	3.05	1.14	14.3	0.1	2.3
J0030+0451	4.87	1.02	4.3	-	8.8
J0340+4130	3.30	0.71	49.6	-	1.7
J0613−0200	3.06	0.96	38.8	1.2	8.6
J0645+5158	8.85	0.49	18.2	-	2.4
J0931−1902	4.64	0.41	41.5	-	0.6
J1012+5307	5.26	1.71	9.0	0.6	9.2
J1024−0719	5.16	1.86	6.5	-	4.0
J1455−3330	7.99	2.43	13.6	76.2	9.2
J1600−3053	3.60	0.95	52.3	14.3	6.0
J1614−2230	3.15	0.96	34.5	8.7	5.1
J1640+2224	3.16	0.28	18.5	175.5	8.9
J1643−1224	4.62	1.85	62.4	147.0	9.0
J1713+0747	4.57	0.85	16.0	67.8	8.8
J1738+0333	5.85	2.41	33.8	0.4	4.0
J1741+1351	3.75	3.02	24.2	16.3	4.2
J1744−1134	4.07	0.89	3.1	-	9.2
J1747−4036	1.65	1.32	153.0	-	1.7
J1832−0836	2.72	0.87	28.2	-	0.6
J1853+1303	4.09	0.87	30.6	115.7	5.6
B1855+09	5.36	1.78	13.3	12.3	8.9
J1903+0327	2.15	1.88	297.6	95.2	4.0
J1909−3744	2.95	1.40	10.4	1.5	9.1
J1910+1256	4.98	0.97	38.1	58.5	8.8
J1918−0642	7.65	2.57	26.6	10.9	9.0
J1923+2515	3.79	0.96	18.9	-	2.2
B1937+21	1.56	10.51	71.0	-	9.1
J1944+0907	5.19	1.73	24.3	-	5.7
J1949+3106	13.14	9.96	164.1	1.9	1.2
B1953+29	6.13	2.97	104.5	117.3	7.2
J2010−1323	5.22	0.48	22.2	-	4.1
J2017+0603	2.90	0.80	23.9	2.2	1.7
J2043+1711	2.38	0.52	20.7	1.5	2.3
J2145−0750	16.05	2.98	9.0	6.8	9.1
J2214+3000	3.12	1.47	22.6	0.4	2.1
J2302+4442	5.19	1.38	13.7	125.9	1.7
J2317+1439	3.45	0.24	21.9	2.5	8.9

Note. — Table adopted from Arzoumanian et al. (2015a).

Kramer 2005) is called the dispersion constant. At Arecibo, multi-frequency observations on the same day are enabled by the agility of the receiver cabin at the Gregorian focus. At the GBT, physical switching between receivers at the prime and Gregorian foci necessitates that there is a gap of a few days between the low and high frequency observations. The asynchronous GBT measurements will introduce stochastic timing errors due to the intrinsically different DMs at the time of each observation; for a characteristic separation of $\sim$3 days between the L-band and 820 MHz observations, this effect is anticipated to be small ($\sim$10 ns, Lam et al. 2015), and will be a non-issue in the era of truly broadband instrumentation.

Table 4.2 lists the various receiver+backend combinations used, their range of dates, and their frequency coverage. All pulsars are observed in the 1.5 GHz range ("L-band"). At the GBT, pulsars are also observed with the 820 MHz receiver, and at Arecibo pulsars are observed either with a higher or lower frequency receiver (i.e., either at "S-band" (2 GHz), or with the 327/430 MHz receivers), depending on their spectral dependence and/or ISM characteristics. The temporal and frequency coverage for each pulsar is shown in Figure 4.1. Two pairs of backend data acquisition systems were used to collect these data, with a few months of overlap during the upgrade: the Astronomical Signal Processor (ASP, at Arecibo) and the Green Bank ASP (GASP, at Green Bank), and the Puerto Rican Ultimate Pulsar Processing Instrument (PUPPI, at Arecibo) and the Green Bank UPPI (GUPPI, at Green Bank). ASP & GASP are essentially clones of one another, as are PUPPI & GUPPI, and both were operated in a mode to record coherently dedispersed, folded profiles.

For the purposes of our analyses and comparisons, we used virtually the exact same set of profiles as the ones from which channelized TOAs were measured and presented in A15. The final data set considered are total intensity profiles averaged to have 20-

Table 4.2. Observing Frequencies and Bandwidths

Telescope Receiver	Plot Color	ASP/GASP Data Span[a]	ASP/GASP Frequency Range[b] (MHz)	ASP/GASP Usable Bandwidth[c] (MHz)	PUPPI/GUPPI Data Span[a]	PUPPI/GUPPI Frequency Range[b] (MHz)	PUPPI/GUPPI Usable Bandwidth[c] (MHz)
Arecibo							
327	red	2005.0−2012.0	315 − 339	34	2012.2−2013.8	302 − 352	50
430	orange	2005.0−2012.3	422 − 442	20	2012.2−2013.8	421 − 445	24
L-wide	blue	2004.9−2012.3	1380 − 1444	64	2012.2−2013.8	1147 − 1765	603
S-wide	purple	2004.9−2012.6	2316 − 2380	64	2012.2−2013.8	1700 − 2404[d]	460
GBT							
Rcvr_800	green	2004.6−2011.0	822 − 866	64	2010.2−2013.8	722 − 919	186
Rcvr1_2	blue	2004.6−2010.8	1386 − 1434	48	2010.2−2013.8	1151 − 1885	642

Note. — The “Plot Color” column serves as a legend for the figure data in this chapter and the appendix of per-pulsar plots. Adopted from Arzoumanian et al. (2015a).

[a]Dates of instrument use. Observation dates of individual pulsars vary; see Figure 4.1.

[b]Most common values; some observations differed. Some frequencies unusable due to radio frequency interference.

[c]Nominal values after excluding narrow subbands with radio frequency interference.

[d]Non-contiguous usable bands at 1700 − 1880 and 2050 − 2404 MHz.

30 min subintervals, 2048 phase bins, and variable resolution for frequency channels: 4 MHz for all ASP & GASP data, and 1.5625/1.5625/3.125/12.5/12.5/12.5 MHz for data from the 327/430/Rcvr_800/Rcvr1_2/L-wide/S-wide receivers with PUPPI & GUPPI. The data considered here are identical with respect to the removal of radio frequency interference (RFI), but some minor differences exist in how some of the L-band GUPPI data were polarization calibrated. Also, there will be negligible differences in a few of the comparisons, since at the time of writing the timing models for a few pulsars[6] are being finalized for the public release of the 9-year data set from A15. These discrepancies should have no effect on our results.

From ASP & GASP to PUPPI & GUPPI

PTA experiments are currently sensitivity-limited; e.g., any improvement in the collecting area, the number of pulsars, receiver sensitivity, or the frequency coverage should improve PTA sensitivity. As stated earlier, NANOGrav is a primary motivation for several surveys to find more pulsars. However, until large "mid-frequency" arrays become the norm, we can expect[7] that 100-meter-class dishes will continue to be the workhorses for pulsar astronomy. Similarly, no significant improvement to receiver sensitivities at our observing frequencies is anticipated.

The one avenue remaining is to expand frequency coverage, and the transition from ASP & GASP to PUPPI & GUPPI represent a first stage of maximizing PTA sensitivity. In the ASP & GASP era, the available processing power limited the observable instantaneous bandwidth, whereas in the PUPPI & GUPPI era, the receiver bandwidth has become a limiting factor (see §1.5). With a $\sim 10\times$ increase of usable bandwidth in some cases, the median TOA uncertainty decreases by a factor of a few

[6] Specifically, J1909−3744, J1944+0907, and J1949+3106.

[7] (though at the time of writing the uncertain futures of both the GBT and AO imply that "hope" may be a better word to use)

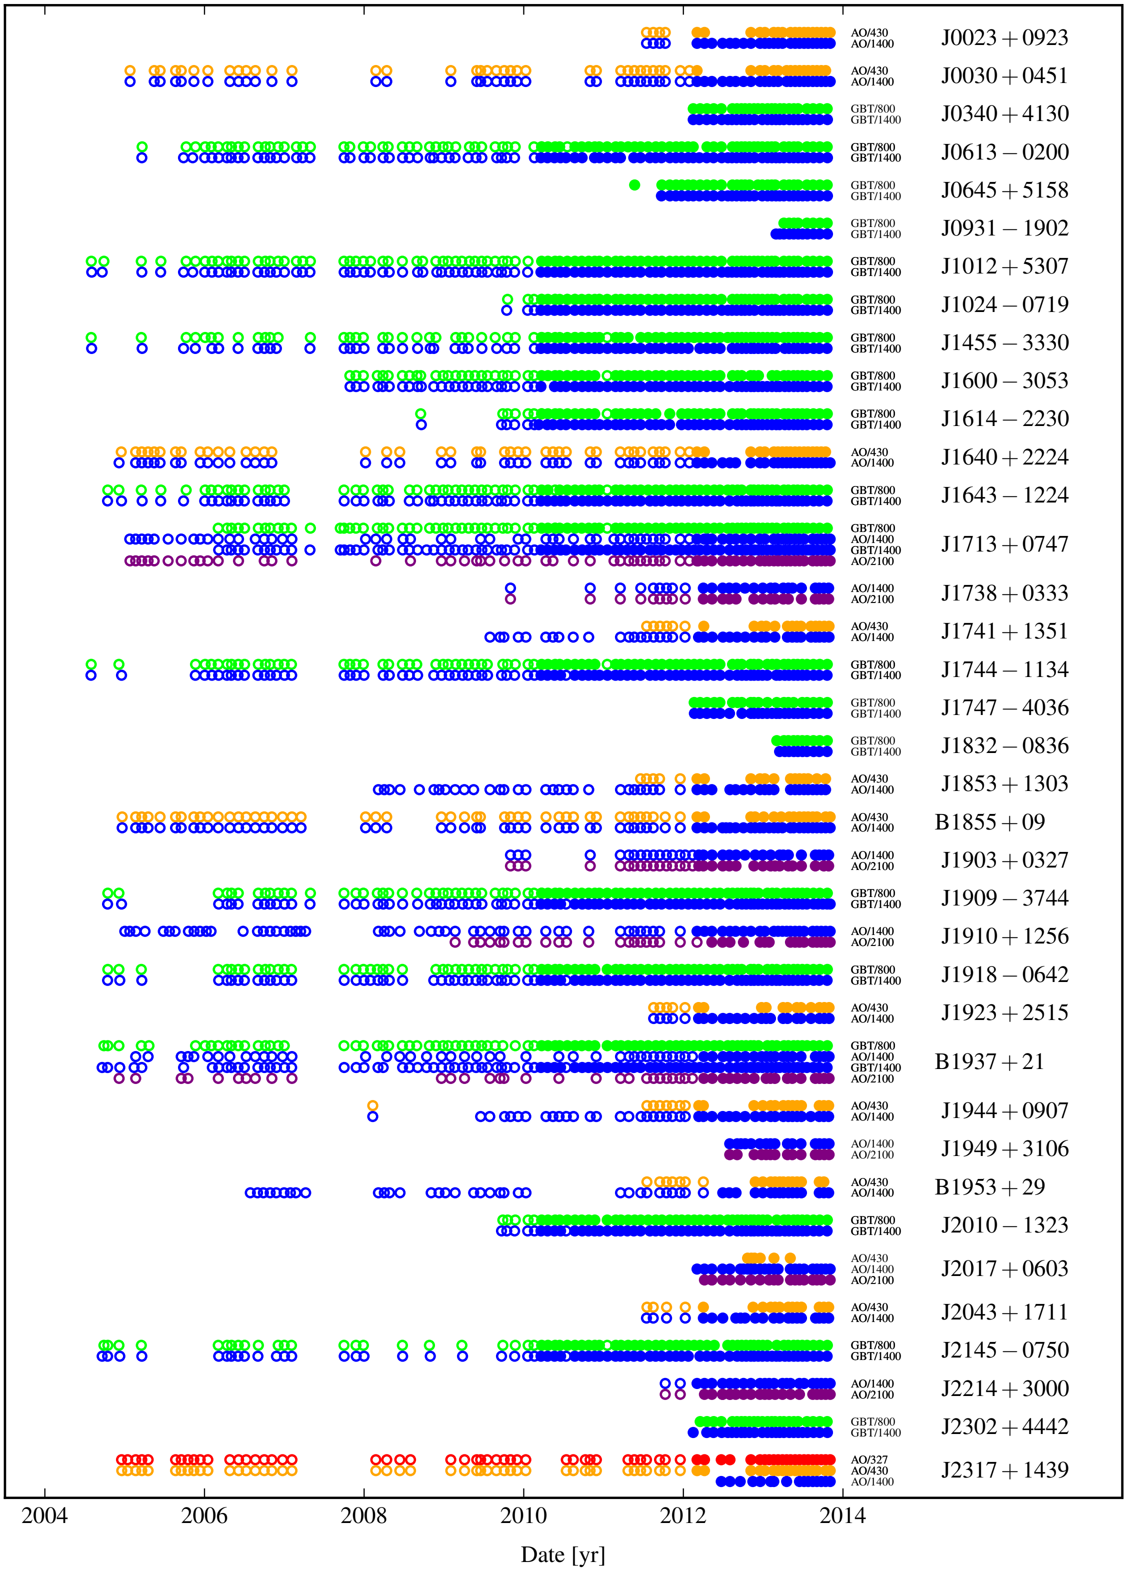

Fig. 4.1 – Observing epochs and frequencies for the pulsars in the NANOGrav 9-year data set. Open symbols are observations with ASP & GASP, closed circles are PUPPI & GUPPI. Figure from Arzoumanian et al. (2015a).

(cf. Table 2 in both Demorest et al. (2013) and A15); a factor of $\sim$3 is expected from the radiometer equation, but this number will be larger or smaller based on the spectrum of the pulsar (which will typically have an index around -1.5) and how large the decorrelation bandwidth for diffractive interstellar scintillation (DISS) is with respect to the bandwidth in question. Examples of the improvement in TOA uncertainty between the two generations of backends can be seen in the timing residual plots at the end of this chapter (§4.4; e.g., Figure 4.89).

As described in §1.5.1 and addressed in Chapter 2, observing a large bandwidth comes with some challenges arising from the fact that one can lose significant timing precision from averaging pulse profiles over more than several tens of MHz at these frequencies. These problems arise from intrinsic pulse profile evolution, as well as pulse broadening from the inhomogeneous ISM, DISS, and having to measure precisely a variable DM. The current solution in the NANOGrav data sets involves channelizing the band into profiles (with the resolutions given earlier) and measuring a TOA per channel using a standard Fourier domain cross-correlation technique (Taylor 1992). However, a single fixed template profile is used for each pulsar for each band, which introduces systematic trends into the residuals as a function of frequency because of the implicit assumption that the profile within a band is independent of frequency.

To account for the profile evolution, NANOGrav did the following. In the 5-year analysis of the ASP & GASP data (Demorest et al. 2013), the profiles were aligned by adding fixed (in time) phase offsets ("JUMPs") between the sets of TOAs at a given frequency. This already undesirable approach would have added almost an order of magnitude more free parameters into the timing model for each pulsar in the entire 9-year data set. Instead, A15 use a set of *ad hoc* parameters to account for the residual structure that remains after fitting the dispersive ν^{-2} law to the channelized frequency TOAs. From A15,

$$\Delta t_{\mathrm{FD}} = \sum_{i=1}^{n} c_i \log\left(\frac{\nu}{1\ \mathrm{GHz}}\right)^i. \tag{4.2}$$

We will refer to the coefficients c_i that describe the excess Frequency Dependence as "FD parameters". FD parameters were included in the timing model if they passed an F-test significance threshold of 0.0027; this criterion was used for all deterministic timing model parameters. The number of FD parameters n needed for each pulsar is given in Table 4.4, which will be discussed in the next section 4.2.2. Of the 37 pulsars, 7 required zero FD parameters, 14 required one, 11 required two, 3 required 3 (B1855+09, J1918−0642, & J2317+1439), 1 required 4 (J1713+0747), and 1 required 5 (B1937+21). To elucidate the effect of using a fixed template to make channelized TOAs in each band, the end of this chapter includes a sequence of frequency residual plots — the larger, well-defined systematics generally require more FD parameters to characterize. These are Figures 4.8–4.44 in §4.4.

Because many of the NANOGrav MSPs are low DM pulsars, several of them scintillate strongly (they suffer from DISS; see Figure 3.4 for a simulated example of DISS). The ASP & GASP bandwidths were relatively small and so DISS often up- or down-weighted all of the TOAs from a given epoch roughly equally. PUPPI & GUPPI, however, cover enough bandwidth to see several scintles. What this means is that, for a given epoch, some of the channelized TOAs will carry significantly more weight than those that were measured from profiles with very low signal-to-noise ratios (S/N), which will have have non-Gaussian uncertainties. These low S/N TOAs are problematic for NANOGrav's timing analysis because the validity of the least-squares approach to fitting a timing model relies on the assumption of input measurements with Gaussian uncertainties. This problem is treated quantitatively in Appendix B of A15, which gives an explicit probability distribution for the TOA uncertainties as a function of S/N.

Currently, NANOGrav mitigates the problem by culling all TOAs that have S/N < 8. This cutoff is drawn in the left subplot of Figure 4.2, which shows the timing residuals from J1455$-$3330 normalized by their uncertainties as a function TOA S/N. A S/N ~ 8 corresponds to approximately where the distribution of normalized residuals approaches a normal Gaussian. In the right subplot, the S/N cutoff corresponds to where the reduced χ^2 from the timing model fit to the TOAs dips below $\sim$2, which equates to cutting 27% of the TOAs. Admittedly, J1455$-$3330 is an extreme example of where DISS plays a significant role, and this 27% of the TOAs carries much less than 27% of the weight of the data, but we'd like to emphasize that our wideband methodology naturally accounts for DISS by appropriately weighting all of the channels in the global fit for a TOA and DM, meaning that all of the information is conserved from the observation and only in cases of an extremely low overall S/N will we have to worry about non-Gaussianities.

To summarize this section, the implementation of PUPPI & GUPPI and the proliferation of backends like them, which are predecessors to truly "ultimate" pulsar instruments, instigated the developments of our wideband methods introduced in two earlier chapters. In what follows, we present the results from applying wideband Gaussian component models to describe profile evolution in the NANOGrav 9-year data set. In a sense, this can be viewed as a natural replacement for FD parameters, JUMPs, S/N TOA cuts, and TOA averaging in the era of broadband pulsar instrumentation. The code that originated in the earlier chapters underwent significant development for this project, and although the entire analysis pipeline itself is not public, the most fundamental pieces of software are all published[8]. Because of the

[8] `https://github.com/pennucci/PulsePortraiture`

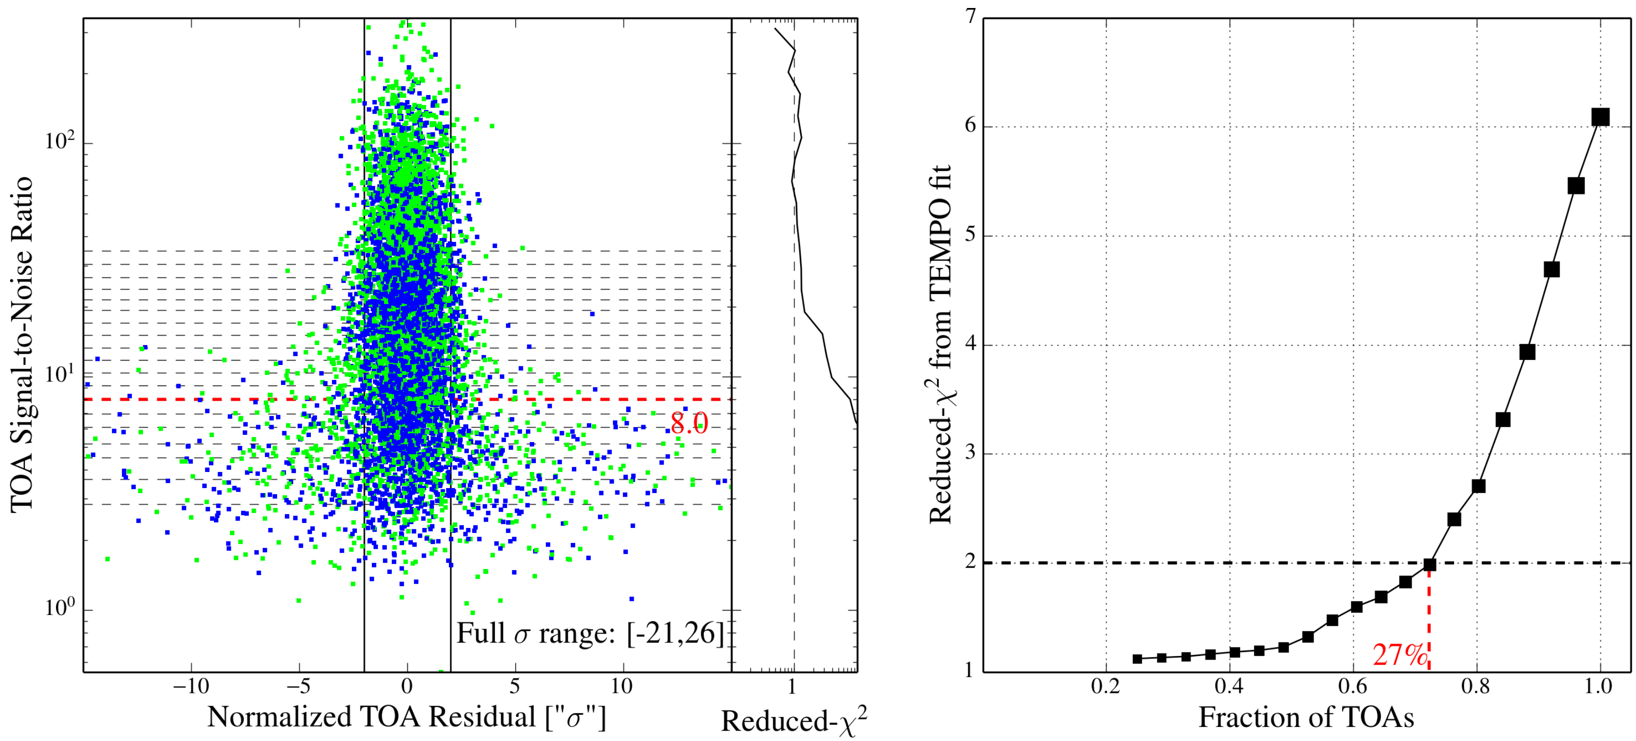

Fig. 4.2 – J1455−3330 has 27% of its TOAs cut in A15 (TOAs with S/N < 8), primarily due to DISS. Our wideband timing method obviates the need to cull any TOAs. *Left:* The horizontal dashed lines demark S/N bins with an equal number of TOAs, and the side panel shows the reduced χ^2 for these bins. *Right:* The reduced χ^2 of the timing model fit as a function of the fraction of TOAs that remain after applying an increasing S/N threshold. The red dashed line corresponds to S/N = 8, and the area of the squares is proportional to the WRMS from the timing model fit.

number of pulsars, many of the "embarrassingly parallel" tasks were accomplished by distributing jobs over a 19+1-node cluster, `nimrod`[9].

4.1.2 Expectations

Although this work aimed to be the first major test of our developments so far, it is important to keep in mind that until data from truly wideband and instantaneous observations are available for PTAs (i.e., data with fractional bandwidth of order unity or greater at gigahertz frequencies), it will be hard to assess the successes and shortcomings of our methodology. In quantitative terms, if one considers the frequency ranges in Table 4.2, the dispersive delays corresponding to 1×10^{-3} cm^{-3} pc across each band (receiver) are: 0.72 μs (S-wide), 1.8 μs (L-wide), 1.9 μs (Rcvr1_2), 3.0 μs (Rcvr_800), 2.5 μs (430), and 12 μs (327). This is to be compared with the delay *between* the band centers (here referenced to Rcvr1_2): -0.81 μs (S-wide), 4.3 μs (Rcvr_800), 20 μs (430), and 37 μs (327). It is obvious that the gaps between the bands provide the majority of the "lever arm" by which one makes precise DM measurements, and so we would be remiss if we had only utilized our in-band DM measurements in our comparisons below (see §4.2.2). Similarly, the amount of profile evolution occurring in these gaps will also be larger, and so complete band coverage with a broadband receiver will allow the DM and profile evolution measurements to be more easily disentangled (Hassall et al. 2012).

With these things in mind, we should only expect large improvements in the cases of large DMs and significant profile evolution (see Chapter 3); results should otherwise be expected to be roughly equivalent.

[9] Nimrod was "a mighty hunter before the Lord".

4.2 Application to NANOGrav MSPs

For each pulsar, the following general protocol was followed:

1. Several high S/N observations from each frequency band were quasi-coherently averaged together (`ppalign.py`).
2. A number of Gaussian component models were fit to the average "portrait", which included varying the number of components, varying the way their parameters evolve, and (in some cases) including a fit for a constant scattering timescale (`ppgauss.py`).
3. TOAs and DMs were measured for all observations and a preliminary timing analysis was performed (`pptoas.py, tempo`[10]).
4. The results were assessed, models were iterated on or fine-tuned, a single model was chosen, and a final set of TOAs and DMs were measured.
5. A noise analysis was performed on the TOAs (`PAL2`[11]), and a final set of timing model parameters and residuals was arrived at.

The next sections detail some of the steps in this protocol, including some of the important and novel developments needed to incorporate the wideband DM measurements. We then compare the results to those from the 9-year data set.

4.2.1 Portrait Modeling

In order to build a high S/N average portrait to which we could fit a reliable Gaussian model, we devised a basic scheme to average the phase-frequency data. This is especially important for the less bright pulsars and/or those which show amplitude modulation with frequency from DISS. For each band, we first initially "incoherently" stacked and averaged several of the highest S/N portraits; this is accomplished with PSRCHIVE's `psradd` routine, which we used with an option to align the data using either the known timing model or a fitted phase offset. This first average portrait was

[10] `http://tempo.sourceforge.net/`

[11] `https://github.com/jellis18/PAL2` (Ellis 2014).

then used to measure a phase and DM in each individual portrait so that they could be "coherently" stacked with respect to the changing DM; the average was weighted in each frequency channel by the channel's S/N. The process was iterated a number of times.

The high S/N average portraits were used as input data to fit for a Gaussian model for profile evolution. The fit included three parameters for each independent Gaussian component, one evolutionary parameter for each Gaussian parameter, a DC offset, a scattering timescale[12], and phase and DM parameters to join/concatenate/align the average portraits from each frequency band. For the fit, the average portraits are normalized to the profile maximum in each channel; removing the spectral shape[13] reduces covariances between amplitude and width parameters in particular, but also highlights the evolution of individual components (e.g., see Figures 4.158 & 4.159).

For each pulsar, a large number of models were tried (at least six). To begin the fit, a 100–200 MHz chunk was identified that had significant signal and profile detail, and was averaged to make a profile. This profile (specifying the model reference frequency $\nu_\circ$) was decomposed by hand in an interactive viewer to get initial Gaussian parameters (location, width, and amplitude) for some number of components, n_{Gauss}. The evolutionary parameters were initialized at zero (no evolution). One of two simple functions was used to describe the evolution of each Gaussian parameter (X) with frequency (ν), either a linear or power-law function:

$$X(\nu|X_\circ, \mathrm{m}_X, \nu_\circ) = X_\circ + \mathrm{m}_X(\nu - \nu_\circ), \tag{4.3}$$

[12]For four pulsars only: J1600−3053, J1643−1224, J1747−4036, & J1903+0327. Scattering was either evident in their profiles or easily separated from fitting a larger number of components. The scattering index was a fixed fit parameter, set to $\alpha = -4.0$.

[13]To be exact, a slightly different function than the spectral shape was removed, since we used the profile maximum.

or

$$X(\nu|X_\circ, \alpha_X, \nu_\circ) = X_\circ \left(\frac{\nu}{\nu_\circ}\right)^{\alpha_X}, \tag{4.4}$$

where the reference and evolutionary parameters for each component $\{X_\circ,\ \mathrm{m}_X,\ \alpha_X\}$ are fitted parameters. Only power-law functions were used for the component amplitudes. Since we normalized the portraits as mentioned, the power-law indices of the components' amplitudes are not spectral indices of flux density[14]. Power-law separations of components and growths in width have been observed and modeled in canonical and millisecond pulsars alike (Thorsett 1991; Xilouris et al. 1996; Kramer et al. 1999; Chen & Wang 2014; Hassall et al. 2012, and references therein.). These references modeled the bulk features over at least one decade in frequency; the frequency dependence of individual sub-components is not well studied in the literature (§15.4 of Lyne & Graham-Smith 2012). Because we are not dealing with even a single decade, a linear model is also warranted as a first-order approximation for evolution. Furthermore, because we decompose the profiles into numerous Gaussian components (which is necessary to characterize the data for precision timing), simple comparisons between what has been previously reported and what we find are not meaningful.

In addition to the four classes of models that include permutations of linear and power-law functions for the location and widths of components, we tried at least two more classes for each pulsar. The easiest conceptual model for profile alignment permits no drifting of the components' locations, and a second one fixes one (or more) components in location to be the "fiducial component" (Hassall et al. 2012). The model residuals did not serve well to discriminate between models; when the models weren't obviously poor choices, the residuals were either contaminated by residual

[14] So long as all components' spectral dependencies are dominated by the spectral index of the average flux, one can approximately recover the components' spectral indices simply by adding the spectral index of the average flux, which would have been a nearly exact recovery if we had instead normalized by the average flux in each channel.

RFI and/or data acquisition artifacts (see §4.2.3), or all had approximately the same χ^2 for the same number of parameters. We relied on quasi-subjective model selection criteria that involved assessing the size of a constant offset in DM as measured in each frequency band (see §4.2.2) and the quality of the timing with respect to the 9-year results. Although a more robust model selection method is highly desirable, we have postponed this aspect of the project until more mature general models are developed and until automated methods exist to generate models (i.e., instead of by hand).

In Table 4.3 we present an overview of the Gaussian model used for each pulsar. The table specifies the number of Gaussian components, the number of fixed/fiducial components, which evolutionary function was used for the locations and widths, and the scattering index at 1 GHz, where applicable. In Figure 4.3, we show frequency histograms for the values of the various evolutionary parameters for all Gaussian components of the final models; the two shades of gray for the linear parameters serve to discriminate between positive (lighter/upper) and negative (darker/lower) slopes. All of the histograms of power-law indices have mean values consistent with zero (no evolution): 0.01(2) for locations, -0.02(9) for width, and 0.10(8) for amplitudes, and have standard deviations of 0.13, 1.2 and 1.4, respectively[15]. For the linear slopes, the averages were similarly consistent with zero: 3(4) μrot MHz^{-1} and -62(49) μrot MHz^{-1}, with standard deviations of about 50 and 610 in the same units[16].

A binomial test of the distributions (the "sign test") yielded only one significant two-tail p-value (of 3×10^{-4}) when assuming that a parameter has equal chance of being positive or negative. By this metric, there were a significantly larger number of negative linear slopes in the width parameter (102) than there were positive ones

[15] Quantities in parentheses are 1-σ uncertainties on the least significant digit.
[16] We use the unit "rot" to denote a rotation, or cycle.

Table 4.3. Summary of Gaussian Models

PSR	n_{Gauss}	n_{Fixed}	Loc. Evol.	Wid. Evol.	$\tau_{1\mathrm{GHz}}$ [μs]
J0023+0923	8	0	POW	LIN	-
J0030+0451	15	15	-	LIN	-
J0340+4130	5	0	LIN	LIN	-
J0613−0200	14	1	LIN	LIN	-
J0645+5158	11	11	-	POW	-
J0931−1902	7	7	-	POW	-
J1012+5307	19	19	-	POW	-
J1024−0719	13	1	LIN	LIN	-
J1455−3330	6	1	LIN	LIN	-
J1600−3053	3	3	-	LIN	9.1(2)
J1614−2230	12	1	LIN	POW	-
J1640+2224	9	9	-	POW	-
J1643−1224	6	0	POW	POW	20.3(4)
J1713+0747	10	0	POW	POW	-
J1738+0333	7	7	-	LIN	-
J1741+1351	8	8	-	LIN	-
J1744−1134	10	10	-	LIN	-
J1747−4036	6	6	-	POW	41.6(6)
J1832−0836	8	8	-	LIN	-
J1853+1303	10	0	LIN	POW	-
B1855+09	11	11	-	LIN	-
J1903+0327	4	1	LIN	LIN	329(4)
J1909−3744	6	6	-	LIN	-
J1910+1256	6	1	LIN	POW	-
J1918−0642	11	11	-	LIN	-
J1923+2515	8	4	LIN	LIN	-
B1937+21	12	0	POW	POW	-
J1944+0907	14	14	-	POW	-
J1949+3106	5	5	-	POW	-
B1953+29	8	8	-	LIN	-
J2010−1323	6	6	-	POW	-
J2017+0603	11	11	-	POW	-
J2043+1711	16	1	LIN	LIN	-
J2145−0750	12	0	LIN	POW	-
J2214+3000	10	1	LIN	LIN	-
J2302+4442	12	12	-	POW	-
J2317+1439	10	0	LIN	POW	-

Note. — A power-law function was always used for the component amplitudes. In most cases, either one (fiducial), zero, or all components were fixed in position. The assumed index to reference the scattering timescale τ to 1 GHz was $\alpha = -4.0$. Quantities in parentheses are 1-σ uncertainties on the least significant digit.

(56). This observation has been documented since almost the discovery of pulsars: low frequency components are, in general, intrinsically broader. It is "classically" explained by an apparent radius-to-frequency mapping geometry (Komesaroff 1970; Cordes 1978), but can also be explained by more modern stream/fan-beam geometries (Dyks & Rudak 2015).

If our interpretation is correct, then it is curious that the next smallest p-value (0.3) does not indicate a significant divide, even though it represents the same test in the distribution of power-law indices of the width parameter, which has a similar total number (77 positive, 90 negative). It is possible that if we filtered out only power-law indices of the width parameter that were significantly positive or negative (i.e., with respect to their measurement uncertainties), then we might find a significant division, but the large covariances of the Gaussian parameters with each other discourage even this simple analysis of the parameter uncertainties. Again, our evolutionary parameters are hard to interpret physically because we are not examining the evolution over even a single decade of frequency and because we are decomposing the complex MSP profiles into many components.

At the end of this chapter, we include a series of plots that show, per pulsar, the concatenated average portraits as aligned in the model fit, the average profiles from each frequency band, the fitted Gaussian model, and the model residuals. See §4.4 for details about the plots.

Number of Gaussian Components

One aspect of model selection that we did examine was how to discriminate between models comprised of different numbers of Gaussian components. Specifically, since in most cases the subjective number of "necessary" Gaussian components can

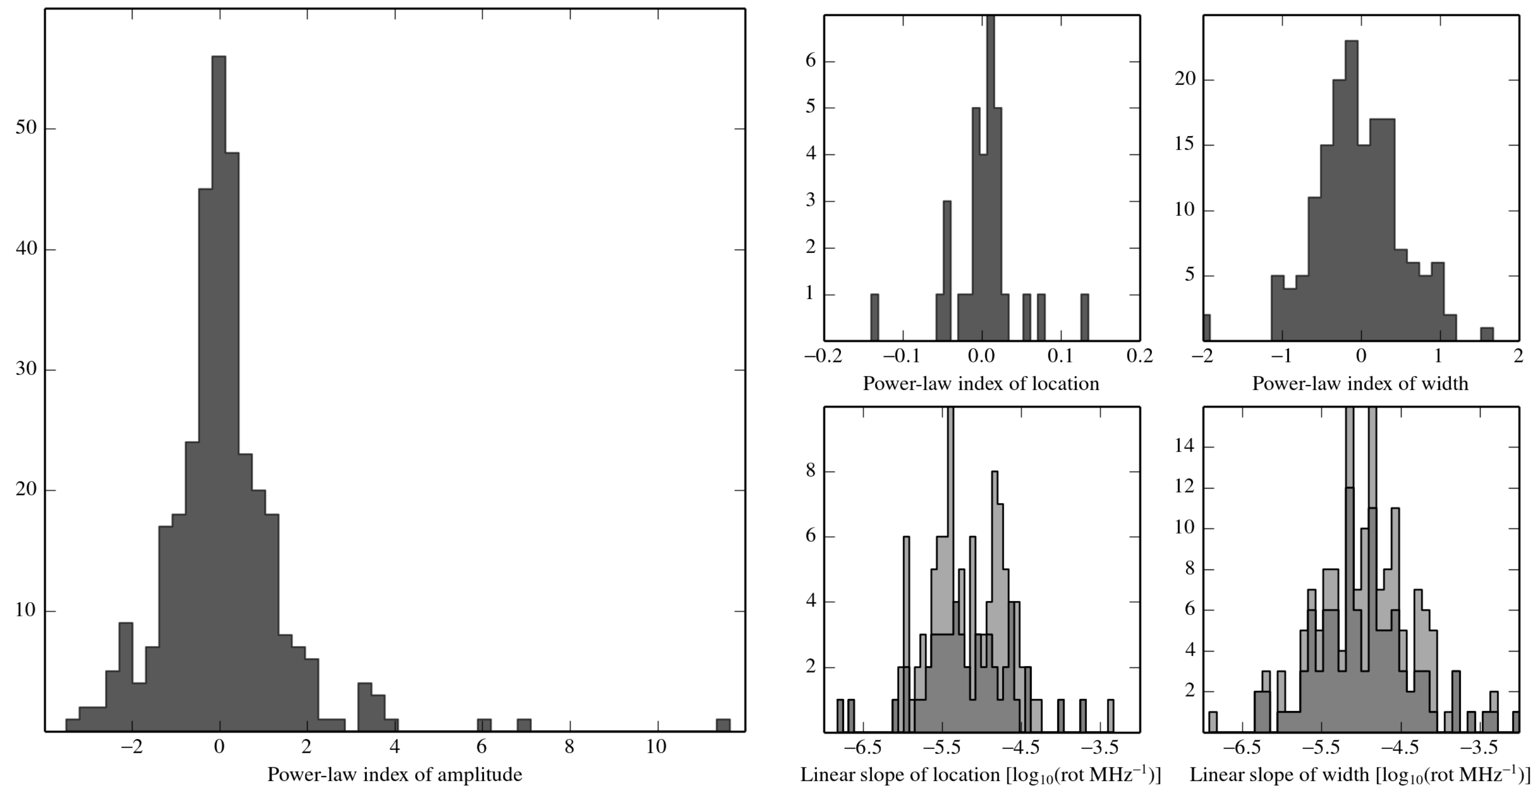

Fig. 4.3 – Distributions of the various Gaussian component evolution parameters. Less than ∼1% of the values are absent in the top-right two plots. The two shades for the linear slope parameters designate positive (lighter/upper) and negative (darker/lower) values. All means are consistent with zero, but the distribution of the linear slope of width parameters has significantly more negative values than positive values (p-value of 0.0003; see text) — this is expected from pulsar emission models.

differ by a few, we were interested in how the number of components affects the TOA measurements. We enlisted a senior undergraduate student, Kevin Stahl, to investigate this question. Kevin looked at several of the NANOGrav pulsars that had various levels of overall brightness and profile complexity and made a variety of models for each. He started with the lowest reasonable number of Gaussian components that could be mapped to the number of profile components and increased this number until the residuals appeared essentially white. One issue with this analysis is that it is difficult to assess the incremental addition of components because there will be large covariances in the components' parameters (e.g., there can be several configurations of three Gaussian components that sum to the same profile component).

Still, the most important result we retained from his analyses was that, as expected, the TOA uncertainties decrease as one increases the number of Gaussians — of course, the better the representation of the model to the data, the fewer remaining systematics. The important takeaway, however, is that in several cases we found that beyond some number of Gaussian components, the improvement in the TOA uncertainty was essentially negligible. Therefore, provided one does not care about an underlying physical interpretation of the components and assuming a critical number of components is reached, having additional or fewer components based on subjective criteria probably has little effect on the precision of the timing. However, because results vary significantly from pulsar to pulsar (unpredictably so), we erred on the side of more components — especially in the cases of very complicated profile shapes and evolution (e.g., J2043+1711, Figures 4.220 & 4.221). These results are summarized in Figure 4.4, which shows the normalized mean of the TOA uncertainty distribution as a function of n_{Gauss} for a subsample of pulsars that have between $\sim$1 and 7 "minimum" components.

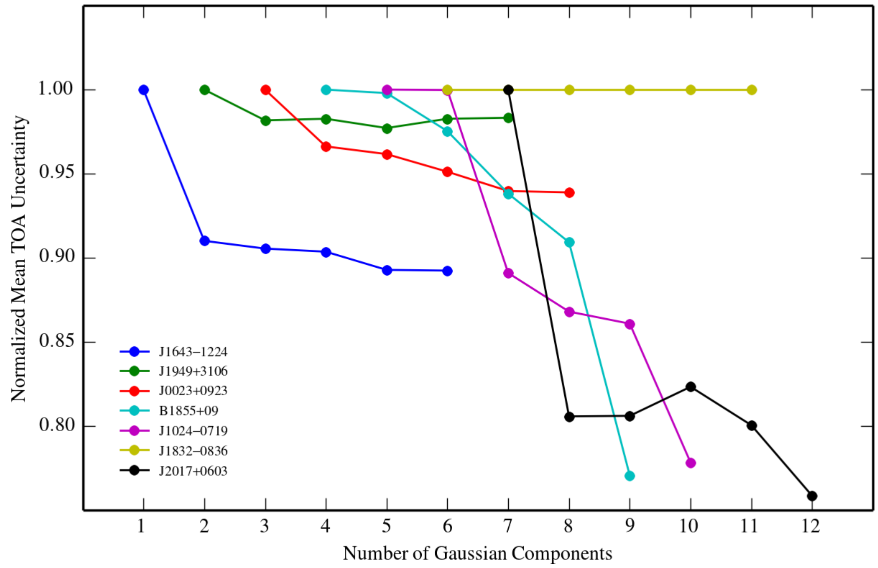

Fig. 4.4 – Analysis of the effect of changing the number of model components. The final models used for all of the pulsars shown either lie in the plateaus of the curves, or comprise more components than shown.

Brief Comments on Specific Pulsars

For the four pulsars from which we have made scattering measurements (J1600−3053, J1643−1224, J1747−4036, & J1903+0327), our values are broadly in agreement with what has been published or predicted[17] (Levin in prep., and references therein.). Within the uncertainties of the measurements and the unknown scattering indices — which can easily differ between ~ -5 and -3 (Lewandowski et al. 2015a) — there are no gross inconsistencies. In the future, we aim to incorporate time-variable wideband scattering measurements separate from the model, which assumes a constant value based on the average portrait. In Chapter 6, we make simultaneous broadband measurements of the scattering timescale and index, as well as the DM, all as a function of time, of the Galactic Center magnetar, J1745−2900.

We tailored our modeling fits for J1713+0747 and B1937+21 because they are both immensely bright (∼10 mJy at 1.5 GHz, which is 3-10 times brighter than the

[17]NE2001 Galactic Free Electron Density Model: `http://www.nrl.navy.mil/rsd/RORF/ne2001/`

others) and are observed in three bands. As will be discussed, the wideband DMs measured in each frequency band for all pulsars were usually offset from one another by a constant. We tried to eliminate the offsets in these two pulsars by choosing to fit the Gaussian model to one bright epoch's multi-frequency data (as opposed to the averaged portrait), so that we could fix the DM as a single parameter across all of the bands, thereby lessening any covariance between the DM alignment parameters and the evolution parameters. In both cases, this approach mitigated the DM offset almost entirely, particularly between L-band and 820 MHz (Figure 4.5). In the future, we can try this method for some of the other significantly bright pulsars.

J1923+2515 and B1953+29 proved to be especially difficult to model using our fairly simple scheme. For J1923+2515, this was because of the specific way that the non-brightest components in each frequency band could incorrectly align with one another based on the position of the brightest component (Figure 4.206). For B1953+29, the disappearance of all but one component in the low frequency band caused an incorrect alignment of the two bright components (Figure 4.214). In both cases, we used the method described above for J1713+0747 and B1937+21, as well as fixing some other parameters by hand[18], to get reasonable models (mostly assessed by eye). For B1953+29, there are residual artifacts in the model because it can be difficult to make components disappear completely, but in both cases we have achieved comparable timing results (§4.2.3).

As a future development to aid in situations like this, instead of extrapolating evolutionary parameters from a single reference profile, a smarter approach might be to "bridge" two very disparate profiles and give better initial parameters.

[18]e.g., note in Table 4.3 that J1923+2515 is the only pulsar with more than one but fewer than all components fixed.

4.2.2 Wideband DM Measurements

The measurements of TOAs and DMs based on a portrait model is covered in Chapter 2. This process was carried out without notable augmentation to the original algorithm. In theory, if profile evolution has been properly modeled and segregated from the dispersive ν^{-2} delay, then the DMs as measured in noncontiguous frequency bands should agree. However, for almost all of our pulsars the wideband DMs measured in different bands displayed constant offsets with respect to one another. The explanation for these offsets is small errors in the model that become apparent when one measures the DM in different frequency bands: whereas the model was fit by linking average portraits from the disparate bands, the measurement of the DM in individual observations from a particular frequency band by fitting a ν^{-2} law will absorb the modeling error, and this error will in general be a function of frequency.

There are two other effects that can cause the DM to be different in the disparate bands. The first arises because the observations are taken at different times[19], and because we are sensitive to DM changes at the level of $\sim 10^{-4}$ cm^{-3} pc, we could in principle see the DM change on short timescales. However, the systematic offsets that we see are of the order 10^{-3} cm^{-3} pc and are hard to explain as intrinsically random-walk DM changes, which are expected to be much smaller in magnitude ($\sim 10^{-6}$ – 10^{-5} cm^{-3} pc) for low DM, weakly-scattered pulsars (Lam et al. 2015; Cordes et al. 2015). Similarly, a frequency dependent dispersion measure could possibly introduce a near-constant offset between DMs measured in different frequency bands, but similar quantitative arguments from Cordes et al. (2015) also rule this out.

Unfortunately, these DM offsets highlight the shortcomings of our modeling. In order to obtain the best set of parameters from the noise modeling (see §4.2.3),

[19] Again, the separation is <1 hour at AO, and several days at the GBT.

we measured and removed the constant offsets on a per-pulsar, per-receiver basis (relative to the DMs measured in L-band)[20]. Their values are plotted in Figure 4.5. Again, a simple interpretation for these numbers is the extent to which the model for profile evolution is in error. From Figure 4.5, one can see that the RMS offset grows between measurements made in the 430 MHz to S-wide receiver bands. For MSPs with spin periods ∼3 ms, the RMS offsets in each band correspond to the same level of modeling error in units of pulse phase bins (∼1–3). In other words, because our modeling routine is limited by the phase resolution of the pulse profile, it can only discriminate between profile evolution and a ν^{-2} law at the level of ∼2 phase bins.

Therefore, as referred to earlier, we are justified in having used the magnitude of these offsets to discriminate between otherwise equivalent portrait models. Our unique treatment of J1713+0747 and B1937+21 discussed above validate our interpretation. In these cases, we confidently are able to segregate the dispersive delay from the profile evolution[21] and identify a fiducial alignment such that the DMs agree between data from both L-band receivers, the 820 MHz receiver, and the S-wide receiver (to within $\lesssim 1 \times 10^{-3}$ cm^{-3} pc ~ 0.5 phase bin).

While acknowledging that we have not modeled profile evolution perfectly[22], we can ask the question of how much better do we do than not modeling profile evolution at all. A rough answer to this question can be reached for each pulsar by comparing the size of the DM offset(s) measured from our best model portraits to the equivalent residual DM that is measured in the post-fit frequency residuals of the channelized TOAs from A15 when no FD parameters are used (Figures 4.8–4.44). These "DM offsets" are obtained by fitting a ν^{-2} dependence to each band's average frequency

[20] In analogy to the constant phase offsets used in `tempo`, this parameter can be thought of as "DMJUMP".

[21] NB: there is always a "normal" covariance between profile evolution and the *absolute* DM.

[22] It is worthwhile to note that the FD parameters only compensate for unmodeled profile evolution since they are coefficients for basis functions of the systematic trends in the frequency residuals.

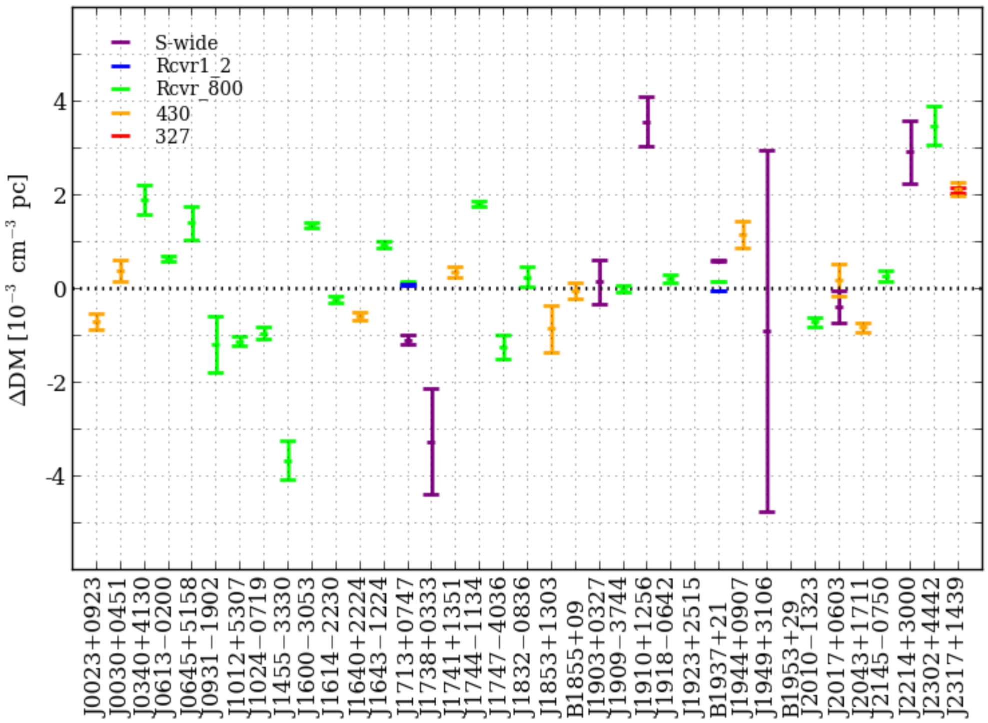

Fig. 4.5 – The measured DM offsets relative to L-band. These arise from imperfect Gaussian modeling of profile evolution. Note the increase in RMS going from low to high receiver bands, which implicates profile resolution as a limiting factor in the modeling. However, the agreement in DMs of J1713+0747 and B1937+21, both of which show the largest levels of profile evolution, validate a modified modeling approach. See text for details.

residuals and are listed in Table 4.4, alongside the DM offset values from Figure 4.5.

22 out of the 37 pulsars showed improvement. Broken up into a dependency on the number of FD parameters used in the 9-year data set to characterize the profile evolution, we see that all but three pulsars requiring more than one FD parameter (16 pulsars) showed improvement. Two of these pulsars require 2 FD parameters, but the values of the DM offsets are comparable (i.e., the profile evolution is equally accounted for). The third pulsar, J2317+1439, which required 3 FD parameters, is unique in that it is the only pulsar with observations at 327 MHz. These data alone are ∼4–5× more constraining of the DM than either the 430 MHz or L-band data. However, the profile at L-band is significantly more detailed than at the lower

Table 4.4. DM Offsets

PSR	# FD	Telescope	DM Offset [10^{-3} cm^{-3} pc]	
	Param.	Receiver	Unmodeled	Gaussian model
J0023+0923	**1**	430	5.75(3)	**-0.71(1)**
J0030+0451	0	430	-0.18(1)	0.38(2)
J0340+4130	**1**	Rcvr_800	4.92(3)	**1.89(3)**
J0613−0200	**2**	Rcvr_800	3.597(17)	**0.628(6)**
J0645+5158	**2**	Rcvr_800	2.27(1)	**1.39(3)**
J0931−1902	0	Rcvr_800	0.83(4)	-1.21(6)
J1012+5307	**1**	Rcvr_800	1.18(2)	**-1.13(1)**
J1024−0719	**2**	Rcvr_800	-2.33(1)	**-0.96(1)**
J1455−3330	1	Rcvr_800	0.77(2)	-3.68(4)
J1600−3053	**2**	Rcvr_800	3.545(8)	**1.339(4)**
J1614−2230	**1**	Rcvr_800	2.482(9)	**-0.244(6)**
J1640+2224	2	430	-0.404(6)	-0.609(8)
J1643−1224	**2**	Rcvr_800	-1.541(26)	**0.924(7)**
J1713+0747	**4**	Rcvr_800	-3.042(8)	**0.139(1)**
		Rcvr1_2	0.3120(77)	**0.0688(8)**
		S-wide	2.839(10)	**-1.104(9)**
J1738+0333	1	S-wide	-1.53(8)	-3.28(11)
J1741+1351	0	430	-0.28(0)	0.34(1)
J1744−1134	2	Rcvr_800	1.118(6)	1.793(5)
J1747−4036	**1**	Rcvr_800	5.02(2)	**-1.25(2)**
J1832−0836	0	Rcvr_800	0.04(2)	0.24(2)
J1853+1303	0	430	-0.20(3)	-0.87(4)
B1855+09	**3**	430	1.04(1)	**-0.06(1)**
J1903+0327	**2**	S-wide	-5.13(3)	**0.13(4)**
J1909−3744	**1**	Rcvr_800	-0.099(1)	**-0.024(7)**
J1910+1256	1	S-wide	1.80(6)	3.55(5)
J1918−0642	**3**	Rcvr_800	-0.376(8)	**0.207(8)**
J1923+2515	1	430	-2.12(3)	10.29(4)
B1937+21	**5**	Rcvr_800	0.1215(119)	**0.1337(2)**
		Rcvr1_2	0.0160(119)	**-0.0462(3)**
		S-wide	1.192(19)	**0.585(1)**
J1944+0907	**2**	430	5.96(4)	**1.15(2)**
J1949+3106	**0**	S-wide	-1.50(42)	**-0.91(38)**
B1953+29	**2**	430	16.69(5)	**12.64(2)**
J2010−1323	**1**	Rcvr_800	-1.910(9)	**-0.720(9)**
J2017+0603	0	430	-0.06(3)	0.17(3)
		S-wide	0.19(2)	-0.41(3)
J2043+1711	1	430	-0.084(7)	-0.837(9)
J2145−0750	**2**	Rcvr_800	5.03(3)	**0.26(1)**
J2214+3000	**1**	S-wide	5.28(9)	**2.90(6)**
J2302+4442	1	Rcvr_800	2.44(3)	3.46(4)
J2317+1439	3	327	-0.299(7)	2.083(5)
		430	-0.22(0)	2.12(1)

Note. — Values of DM offsets from Figure 4.5 compared with analogous DM offsets in the frequency residuals without profile evolution modeling (Figures 4.8–4.44). Bold values are "improvements". Note the distribution of required FD parameters from A15: 1 pulsar with 5, 1 with 4, 3 with 3, 11 with 2, 14 with 1, 7 with 0, for a total of 54 FD parameters. There are 43 "DMJUMPs".

frequencies (see Figure 4.228), and so this is where we referenced our model ($\nu_\circ =$ 1300 MHz). Consequently, the evolution at the distant 327 and 430 MHz bands will be poorly constrained. Note that in both cases the DM offsets are the same for the 327 and 430 MHz data. J2317+1439 was also the pulsar that showed the largest improvement in the timing residuals after removing the offsets from the wideband DMs, again indicating the influence of the 327 MHz data.

For the pulsars that required 1 FD parameter, 8 of 14 showed improvement, and only 1 of the 11 pulsars with no "significant" pulsar evolution (0 FD parameters) showed improvement — and not very significantly so. In a number of these cases where no improvement was seen, the DM offset values are not very different[23]. These results indicate that in a small number of cases we may have artificially introduced profile evolution when there was none. The models for the two pulsars with the largest offsets, J1923+2515 and B1953+29 (the latter of which shows "an improvement"), we have already discussed earlier.

The wideband DM measurements are shown in the middle panels of the timing summaries at the end of this chapter, (§4.4, between Figures 4.45 & 4.117); they are colored by receiver, as with the other plots. Again, the offsets discussed above have been removed. The wideband DMs measured from the ASP & GASP data are included for completeness, but they are in general non-informative due to the small bandwidths and have been plotted without error bars for clarity.

Wideband DMs & DMX

The current model for DM(t) used in the NANOGrav analyses is a piecewise constant function (Demorest et al. 2013; Arzoumanian et al. 2015a). The width of

[23]The uncertainty measurements in Table 4.4 may not be entirely meaningful, since we don't address the goodness-of-fit.

an epoch that determines the DM for each "piece" is usually set to be no bigger than 14 days, with a small number of exceptions when there were only single-frequency data available (usually from the ASP or GASP era), but see A15 for details. In A15, all of the channelized TOAs that fall within one of these DM epochs are fit for a single DM; all such DMs are determined simultaneously with other timing model parameters (including the FD parameters) in a generalized least-squares (GLS) fit using `tempo`. This model is referred to as "DMX".

As mentioned earlier, the frequency gaps between bands carry as much or significantly more dispersive delay than within the bands themselves. This implies that we must also make use of the measured delay between our wideband TOAs to measure each epoch's DM (i.e., we still need to use DMX). The difference is that, instead of having channelized TOAs and FD parameters, we have only 2 or 3 TOAs per DMX bin, which have, in principle, already accounted for profile evolution and measured (in-band) DMs. In order to make full use of the same information, we also need to incorporate our wideband DMs as data into the timing model fit. This was accomplished by augmenting `tempo`'s GLS χ^2 calculation to also include input DMs for each epoch[24]. Schematically, with reference to Equation 1.2,

$$\chi^2_{\mathrm{total}} = \chi^2_{\mathrm{TOA}} + \chi^2_{\mathrm{DM}}, \tag{4.5}$$

where

$$\chi^2_{\mathrm{DM}} = \sum_{i}^{\#\ \mathrm{DMs}} \left(\frac{\mathrm{DM}_i - (\overline{\mathrm{DM}} + \mathrm{DMX}_j)}{\sigma_{\mathrm{DM},i}}\right)^2. \tag{4.6}$$

Here, $\overline{\mathrm{DM}}$ is the mean dispersion measure, DMX_j are the differences from the mean, and j indexes the epoch of the DMX bin in which DM_i should fall ($j \leq i$). The

[24]Thanks to P. Demorest.

overall χ^2_{total} is minimized accordingly[25].

Because the input DM_i provide a separate set of constraints on $\mathrm{DM}(t)$, the average DM must also be fit in `tempo`. This is not usually the case when implementing DMX, which fits for DM differences. Again, the overall profile evolution and average absolute DM are covariant, so it is not meaningful to compare the values of $\overline{\mathrm{DM}}$. However, it is useful to see how the fitted DMX values compare.

The comparisons between the DMX values from our wideband measurements and A15 are shown in subplots as part of the timing comparison figures at the end of this chapter (§4.4, between Figures 4.46 & 4.118). The same DMX epochs were used for all measurements. There is a remarkable agreement in most pulsars, verifying in part that our wideband TOAs are credible. That the DMX trends are much smoother and are more precisely measured than the PUPPI & GUPPI wideband DMs again highlights the importance of the dispersive delay *between* the bands.

In principle, instead of measuring in-band DMs and feeding them into the augmented DMX fit with `tempo`, we could have performed a more analogous set of wideband DM measurements, pairing the low and high frequency portraits together and measuring a single DM across both. This, however, would introduce correlations into the phase (TOA) measurements and would not eliminate systematic residuals resulting from the fit of a single DM or inaccurate modeling. Since either approach should yield nearly the same result, we proceeded as described above. We made this choice because we have developed our methods for use with truly wideband receivers and so we don't find it necessary here to make additional developments to accommodate the lack of simultaneous, wideband observations.

[25] As of June 2015, to enable this functionality in `tempo` (version 13.000) one must use the GLS flag ("-G") and have "DMDATA 1" in the ephemeris file. `tempo` will then make use of DM measurements associated with each TOA specified by the TOA flags "`--pp_dm` #" (the DM) and "`--pp_dme` #" (its 1-σ uncertainty), both of which are outputted by default by `pptoas.py`.

4.2.3 Timing Results

We used the timing models from A15 as initial parameters for our timing analyses. We fit the same astrometric, spin, binary, DMX, and "JUMP" parameters; the FD parameters, of course, were removed from the fit. JUMP parameters are still required to account for the unknown differential latencies between the various receiver setups and the backend data acquisition system[26]. JUMPs are not needed between the two pairs of backend instruments because of a novel method used to independently measure the offsets, which are held fixed and added to the TOAs (Arzoumanian et al. 2015a).

A generalized least-squares (GLS) approach to the timing model fit is required because simpler assumptions about the input data would be in error; TOAs can be time-correlated and their uncertainties, aside from being inhomogeneous, may inaccurately represent the noise in the timing residuals. It is critically important to robustly characterize the noise for high precision timing experiments. The particular noise analysis used here and in A15 is a version of what has been developed in recent years specifically for the cause of PTA experiments (van Haasteren & Levin 2013; Ellis et al. 2013; Ellis 2013; van Haasteren & Vallisneri 2014, 2015; Arzoumanian et al. 2014; Ellis 2014). §5 and Appendix C of A15 give details of this noise modeling scheme, an overview of which will be presented next, followed by our final timing results.

[26]NB: These constant time offsets are covariant with the absolute DM measurement, but will be unecessary when a single broadband receiver is used, or if we can independently measure the delays.

Noise Analysis

TOAs that are uncorrelated in time and frequency may have inaccurate uncertainties, which are nominally estimated from the template-matching process[27]. The noise model accounts for additional sources of "white noise" by transforming the estimated TOA uncertainties $\hat{\sigma}_{i,k}$ as

$$\sigma_{i,k} = E_k \, (\hat{\sigma}_{i,k}^2 + Q_k^2)^{1/2}, \tag{4.7}$$

where i indexes the TOA and k indexes the receiver+backend setup. E_k and Q_k are called "EFAC" and "EQUAD", respectively, in `tempo`. In the absence of EQUAD, EFAC can be thought of in terms of the "usual" scaling by the reduced χ^2 value. In principle, EQUAD and EFAC can account for, e.g., systematics left over in the template-matching procedure and/or small deviations from Gaussian statistics.

A second noise term used in A15 accounts for TOAs that are 100% correlated in frequency, but completely uncorrelated in time, called "ECORR". One physical culprit for non-zero ECORR is intrinsic pulse phase jitter (Cordes & Shannon 2010). Pulse phase jitter is a measure of how much individual pulses jump around in phase due to processes at the neutron star (hence the frequency correlation) that are uncorrelated in time. Therefore, the only way to reduce intrinsic pulse phase jitter is to integrate longer per TOA. We do not model ECORR in our wideband TOAs since we only have one measurement per frequency band. Any source of noise that would be modeled by ECORR in the channelized TOAs would be absorbed by EQUAD.

Finally, time-correlated signals in the residuals are modeled as a red noise Gaussian

[27]For the wideband TOAs, the template is the two-dimensional model portrait; for the channelized TOAs, it is the band-constant model profile.

process whose power spectrum P is a power-law,

$$P(f) = A_{\rm red}^2 \left(\frac{f}{f_{\rm yr}}\right)^{\gamma_{\rm red}}, \tag{4.8}$$

where $A_{\rm red}$ [μs yr$^{1/2}$] is the amplitude of the red noise, $\gamma_{\rm red}$ is its spectral index, and $f_{\rm yr} = 1$ yr^{-1}. Achromatic red noise can be a manifestation of intrinsic "spin noise", which can be thought of as how much the clock drifts over long timescales and will have a fairly steep index ($\gamma_{\rm red} \sim -5$ (Shannon & Cordes 2010)). Spin noise is usually much larger in young pulsars, canonical pulsars, and magnetars — all of which have larger magnetic fields and longer periods. Spin noise ultimately eliminates the possibility of including hundreds of bright, slow pulsars into a PTA, although investigations on this front are of current interest to some. Chromatic red noise will arise in the residuals as a result of inaccurate/incomplete modeling of ISM effects and will have a shallower noise spectrum. For instance, a time-variable scattering timescale will cause measurable changes in the shape of pulse broadening. Consequently, the dispersive delay will be measured and removed inaccurately, leading to chromatic structure (for instance, see the discussion of J1643−1224 that follows).

The estimation of the noise parameters is done by Markov Chain Monte Carlo sampling of the joint posterior distribution, where the process is greatly reduced in complexity by first analytically marginalizing over the timing model parameters. We found that there are significant covariances between EFAC and EQUAD in particular, but retained the same protocol of using the maximum likelihood values in the timing analysis. In addition to our EQUAD values absorbing ECORR, the EQUAD values from A15 may be susceptible to contamination by unmitigated RFI in specific channels or particular artifacts that arise from imperfect acquisition of the data[28].

[28]PUPPI & GUPPI sample the time series data using interwoven samplers, which are not perfectly

In principle, wideband TOAs should be less vulnerable to these effects, although this claim has not yet been investigated.

Red Noise Comparison

For the above reasons, it is not meaningful to directly compare the two sets of EFAC and EQUAD parameters. However, more important are the red noise parameters, which one expects to be similar. Here, we only consider the ten MSPs whose Bayesian noise modeling favored inclusion of the red noise parameters in A15. The red noise parameters for a pulsar were included in its timing model if the Bayes factor exceeded one hundred[29]. Bayesian model selection is very computationally expensive (compared to parameter estimation); at the time of writing we do not have Bayes factors for our noise models, and so cannot assess the differences in significance of the red noise. We will simply assume that the values we measure for these ten pulsars are similarly significant, making note that their posterior probability distributions were significantly peaked.

In the left subplot of Figure 4.6 we plot our spectral indices of the red noise, γ_{red}, along with those from A15 and those from a recent publication by the European Pulsar Timing Array (EPTA; Lentati et al. 2015b). The dashed and dotted lines are the prediction from Shannon & Cordes (2010) for γ_{red} ($= -5 \pm 0.4$) in the post-fit residuals[30] based on a collection of canonical and millisecond pulsars. As is noted in

---

"180°" out of phase with one another, but are assumed to be. This results in an effectively different sampling rate and causes aliasing of the signal back into the band. The "ghost" of the pulsar signal is band-reversed and so it usually is smeared out by dedispersion and/or lies below the noise level. However, for certain bright pulsars with the "wrong combination" of spin period and DM, the artifact will bias TOAs from frequency channels near to where the ghost and primary signals cross. This effect explains the outlier points in, e.g., Figure 4.24 and the diagonal signal seen in e.g., Figure 4.189.

[29]The Bayes factor here is the ratio of the total probability of observing the timing residuals including the red noise model (marginalized over all possible parameters), to the total probability of observing the residuals without the red noise model.

[30]By "post-fit" we mean after the deterministic timing model has been subtracted, where the

A15, most of the pulsars (here six, instead of seven, out of ten) are inconsistent with the prediction.

In the right subplot of Figure 4.6 we show the estimated RMS amplitude of the post-fit timing noise in our data ($\sigma_{\mathrm{TN},2}$) as a function of the prediction from Shannon & Cordes (2010) ($\sigma_{\mathrm{TN},2}^{\mathrm{SC10}}$). The same subplot in Figure 3 of A15 is nearly identical with respect to the ten MSPs with detected red noise. Hence, we conclude that we measured the same amplitude of red noise in these ten MSPs[31]. The gray points in the figure are the 95% upper limits on $\sigma_{\mathrm{TN},2}$ for the other 27 MSPs, many of which lie below the line of equality. As is also concluded in A15, our red noise amplitudes are inconsistent with the prediction of Shannon & Cordes (2010).

As explained in A15, many of the shallow values for γ_{red} may be a result of ISM effects or incomplete modeling of DM(t), as opposed to intrinsic timing instability. The only notable differences in γ_{red} from our analyses are for J0613−0200, J1903+0327, and B1937+21.

Working in reverse, B1937+21's red noise index is marginally consistent with that from A15, being slightly more negative. Looking at both sets of residuals and DMX curves (i.e., Figure 4.97 and specifically the *averaged*, non-whitened residuals from A15), we see that the wideband results are smoother at several epochs where there are small discontinuities or high frequency trends in the A15 results. In particular, we are pointing at the correlated changes in the DMX values and the residuals near 2011.8 and 2012.2, where the transition from ASP/GUPPI to PUPPI/GUPPI occurs. The removal of these high frequency features could steepen the red noise, closer to the value observed in an analysis of this pulsar that covered a time span 2.5 times longer (Shannon & Cordes 2010).

relevant component to the discussion here is the long-term spin down parameter, $\dot{P}$, which effectively removes a quadratic.

[31] J1903+0327's red noise is less constrained, but its maximum posterior amplitude is the same.

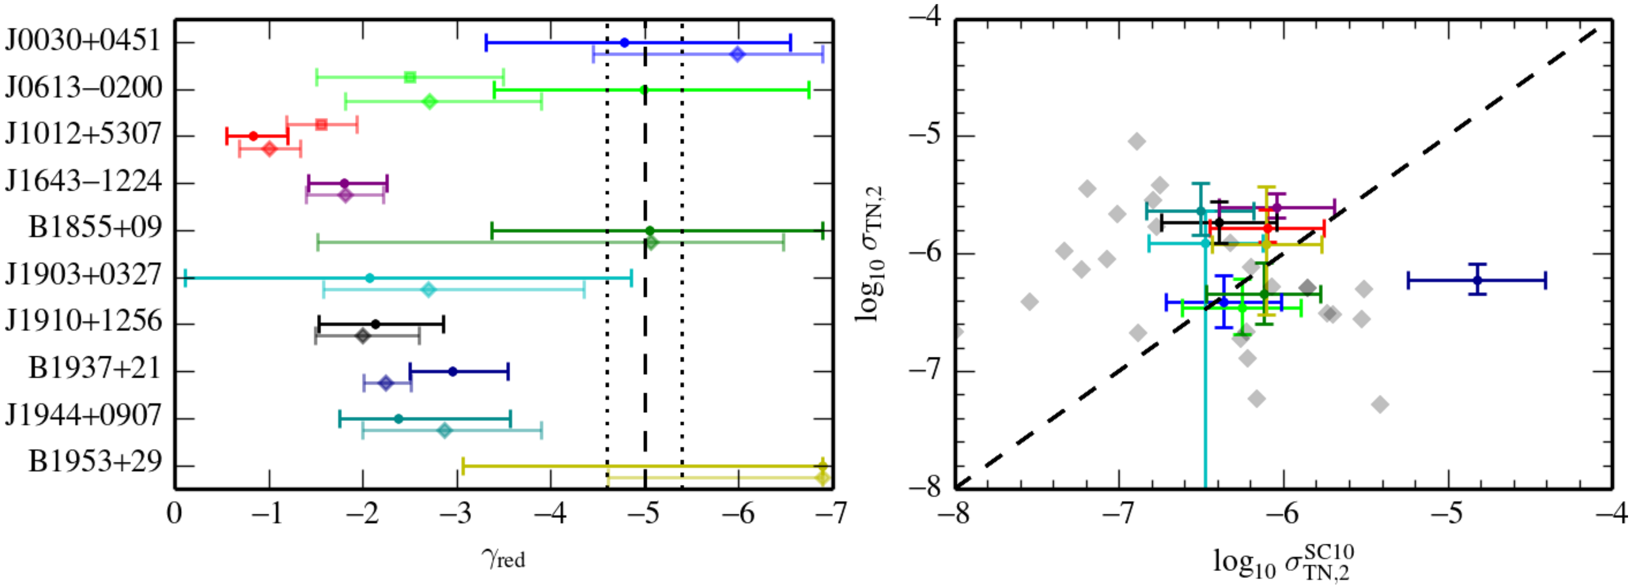

Fig. 4.6 – Comparison of the red noise parameters. *Left:* The measured red noise power-law index from the power spectrum of the post-fit residuals, $\gamma_{\rm red}$. This parameter is shown for the ten pulsars that exceeded the Bayes factor criterion (>100). The dots are from the noise analyses of the wideband TOAs, the diamonds are reproduced from the analyses of the channelized TOAs from A15, and the squares are from Lentati et al. (2015b). The dashed and dotted lines are the prediction from Shannon & Cordes (2010). *Right:* The measured and predicted values of the RMS amplitude of the post-fit red noise. The points are colored the same as in the left subplot; the gray diamonds are the 95% upper limits for this quantity from the other 27 pulsars. All points with error bars are shown with maximum *a posteriori* parameter values and 68% credible intervals.

Our value of $\gamma_{\rm red}$ for J1903+0327 agrees with that from A15, but is much less constrained. J1903+0327's wideband residuals show much more high frequency structure than in A15, although the overall amplitude appears to be somewhat less (and also less constrained). J1903+0327 displays the largest amount of pulse broadening from scattering (Figure 4.198) and is the highest DM pulsar in the sample ($\sim$300 cm^{-3} pc). These facts may combine with our profile evolution model (which included a constant scattering timescale of $\sim$330 μs at 1 GHz) to introduce achromatic high frequency features.

J0613−0200's red noise index is the only significantly different value, now being consistent with the predicted −5. We can only note that we see a smoother cubic-like trend in our residuals over the full 9-year span than is shown in A15. We also note that we have modeled its complex profile evolution well (Figures 4.162 & 4.163; its DM offset relative to L-band is $\sim$1 phase bin).

J1643−1224: Variable Scattering or Chromatic DM?

One other pulsar warrants some discussion. J1643−1224 shows significant, shallow red noise in both sets of analyses. However, whereas in A15 the red noise appears highly chromatic, we see no such frequency dependence in the residuals (Figure 4.69). The chromatic trends arise because the *apparent* DM in each band is different enough so that when a single DM is measured across both bands with DMX, the correction to infinite frequency of the channelized TOAs is sufficiently inaccurate to be evident in the residuals. This explanation is supported by the strong annual correlations between the DMX trend and the 820 MHz residuals from A15. The apparently chromatic DM is manifested by pulse broadening from interstellar scattering and/or an intrinsically frequency dependent dispersion measure, DM(ν). The effect of a constant scattering time or a constant departure from the usual frequency independent DM would be absorbed in the FD parameters and/or the DMX values, but time variability of either effect would induce chromatic trends. Cordes et al. (2015) predict that the form of DM(ν) should change on a refractive timescale and J1643−1224 is known to have previously exhibited unusual scattering behavior (Maitia et al. 2003), so either effect may be at work.

The reason why the wideband TOAs do not show a chromatic dependence is explained by the fact that there are only two points in each DMX bin; to minimize the WRMS residual, DMX assigns the DM offset corresponding to a ν^{-2} law between the points[32]. In comparison, the unmodeled chromatic effect (whatever it is) is revealed in our wideband DM measurements, which vary significantly on timescales shorter than a year and are, at times, anti-correlated between the DMs measured in either band. In Figure 4.70, we see that the DMX values agree with those from A15 (with a slight offset for epochs after 2013), implying that the large scale variations in the DM — the overall linear trend from relative motions and the annual variations from the line of sight oscillating over the inhomogeneous ISM — are being mitigated almost identically. However, our TOAs are more accurate than the averaged TOAs from A15. That is, the red noise trend is not annually, chromatically variable. This is true because our "TOA averaging" happens at the time of the simultaneous TOA and DM measurement as opposed to after an incorrect dispersive delay across both bands is removed from the channelized TOAs. Another way to look at this is that the apparent DM is different in each band, but is separately mitigated only in the wideband measurements. We also note that our TOA residuals and DMX curves are essentially unchanged if we disallow the wideband DMs from being considered in the

[32]Note, however, that this fit is done simultaneously between all DMX values as well as with the other timing model parameters.

timing model fit[33] and we have verified that, in this case, the average wideband DM values track the DMX values closely.

It is difficult to say which, if not both, effects are at work. Two initial observations seem to implicate a frequency dependent DM over a variable scattering timescale. First, the EFAC parameters for the channelized GUPPI TOAs are very close to 1.0 (within 10%), which is what you expect if the profile template matching is accurate; i.e., pulse shape changes are not the culprit. Second, in the case of variable scattering one naively expects the two wideband DM trends to always be correlated. Instead, we see periods of anti-correlation, which can happen in the case of a frequency dependent DM (Cordes et al. 2015). Neither of these observations is convincing, but we next quantify the expected RMS of the wideband DMs based on predictions from Cordes et al. (2015) for a frequency dependent DM.

Using our estimate for the scattering timescale at 1 GHz ($\sim$20 μs, in agreement with others, e.g., G. Jones, unpublished), the expected RMS difference in DMs measured at L-band and 820 MHz ($\sigma_{\overline{\mathrm{DM}}}$) is $\sim 1 - 3 \times 10^{-3}$ cm^{-3} pc, which agrees well with what is seen in the wideband DMs. Translating this into an RMS timing error in the residuals, $\sigma_{t_\infty,\delta\overline{\mathrm{DM}}} \sim 3 - 4$ μs, which is not very different from the 2 μs WRMS residual reported in A15. Furthermore, based on the distance estimate to J1643−1224 of 4.9 kpc (Toscano et al. 1999), we can derive a reasonable transverse velocity of the pulsar across the line of sight through the ISM, $V_{\mathrm{ISS}} \sim 240$ km s^{-1} (§4.2.6 of Lorimer & Kramer 2005). This allows us to estimate the timescale for refractive scintillation (Δt_{RISS}) over which time the form of DM(ν) will change (Cordes et al. 2015). For J1643−1224, $\Delta t_{\mathrm{RISS}} \sim$ month, which is in line with the timescale of the wideband DM oscillations.

We conclude that a frequency dependent DM may be responsible for the chromatic

[33] i.e., "DMDATA 0"

red noise in the residuals of J1643−1224 from A15 and the chromatic wideband DMs. The former is mitigated by having measured the DM and TOA in each band independently (leading, in fact, to the chromatic DMs). We should also note that J1643−1224 is known to lie behind the HII region ζ-Ophiuchus, which is likely to blame for all chromatic phenomenology seen here (Arzoumanian et al. 2015b). It will be interesting to test these hypotheses in future iterations of the modeling and measurements via simulations and supplementary observations.

If our interpretation of J1643−1224's results is correct, a natural question to ask is why a few other pulsars still show chromatic red noise features in their residuals — particularly J1600−3053 and J1747−4036, which are singled out in A15. Both sources are high DM pulsars ($\sim$50 and $\sim$150 cm^{-3} pc, respectively) for which we have included scattering as part of their model portraits ($\sim$10 and 40 μs at 1 GHz, respectively). In both cases, the chromatic trends are not nearly as pronounced (they are essentially absent in the wideband analysis of J1747−4036), and the wideband DMs are not as precisely measured as in J1643−1224. For J1600−3053, we can again estimate $\sigma_{\overline{\rm DM}}$ from our scattering measurement, which we find to be on the order of $\sim 0.5 \times 10^{-3}$ cm^{-3} pc, or roughly the size of the wideband DM uncertainties for this pulsar. A similar story is told for J1747−4036, which may have $\sigma_{\overline{\rm DM}} \sim 1 \times 10^{-3}$ cm^{-3} pc, but its wideband DM uncertainties are also as large as this. Lastly, J1903+0327 is the largest DM ($\sim$300 cm^{-3} pc), most heavily scattered ($\sim$330 μs at 1 GHz) pulsar in the sample, but shows no prominent chromatic trends. This may be a consequence of choosing S-band as the second frequency band in which to observe this pulsar. However, the predicted level of $\sigma_{\overline{\rm DM}}$ for this pulsar is still large, $\sim 5 \times 10^{-3}$ cm^{-3} pc, corresponding to $\sigma_{t_\infty, \delta\overline{\rm DM}} \sim 4$ μs. While we do see inconsistencies in the wideband DMs at the level of several $\times 10^{-3}$ cm^{-3} pc, it is difficult to reconcile the observations with our expectations of DM(ν).

Timing Comparison

At the end of this chapter, in Figures 4.45–4.118 we show all timing summary and comparison plots from our analyses on a per-pulsar basis. The timing summary plots have panels containing our wideband residuals, "whitened" residuals (for pulsars with red noise), wideband DM measurements, and DMX trends. For the red noise pulsars, the residuals are "whitened" by subtracting the maximum likelihood realization of the red noise. In the comparison figures, we plot our wideband residuals versus the averaged residuals from A15[34] in the left subplot, and our DMX values versus those from A15 in the right subplot.

In general, there is reasonably good agreement across all pulsars. A few long-term-timed benchmark pulsars are worth pointing out for some combination of being bright, having red/white noise, having simple/complicated profile evolution, and/or having a substantial trend in DM(t): J1713+0747, B1855+09, J1909−3744, B1937+21, J2317+1439. Because of the large covariances between EFAC and EQUAD, we cannot trust all of the timing results until a more scrupulous method for extracting the best noise parameters is devised.

To consolidate the comparisons of the timing measurements and to highlight the potential influence of the EFAC and EQUAD parameters, we made histograms of the information contained in the comparison plots. Figure 4.7 demonstrates how, on average, the wideband measurements and their uncertainties compare with their counterparts from A15. The bottom row shows the histograms of the differences between the wideband and channelized-TOA measurements, normalized by the square-root of the quadrature sum of their uncertainties. The top row simply shows the histograms

[34] These will be the whitened residuals in both cases, where applicable.

of the ratio of the uncertainties (wideband/A15) of the residuals and DMs from the comparison plots.

For the overall agreement, the histogram of normalized residual differences has a mean of $(-4 \pm 300) \times 10^{-3}$ and a standard deviation of 0.5, and the histogram of normalized DMX differences has a mean of $(-6 \pm 90) \times 10^{-2}$ with a standard deviation of $\sim$1. For the overall ratio of uncertainties, the histogram of TOA uncertainty ratios has a mean of $1.4^{+0.1}_{-0.6}$ with a standard deviation of $\sim$2, and the histogram of DMX uncertainty ratios has a mean of $1.5^{+0.6}_{-0.6}$ with a standard deviation of $\sim$1. The intervals quoted here contain the 16–84 percentile range. Although we see that typically the results agree and we do no better or worse in our measurements, there are some fat tails in the distributions, particularly in the ratio of the uncertainties. We attribute these to differences in the noise modeling.

With these caveats in mind, we next present comparisons of the WRMS timing residuals and estimated timing model parameters. Table 4.5 contains the results from our GLS timing model fits with `tempo`, which was augmented with "DMDATA" to make use of the "DMJUMP"-corrected wideband DMs (see §4.2.2). The total number of degrees of freedom (which includes both the number of wideband TOAs and DMs) is given for each pulsar. Despite the noise modeling described earlier, all pulsars had a reduced χ^2 greater than 1.0. The reason for this is that the noise modeling does not yet incorporate similar noise terms for the DM data points, and our wideband DMs appear to have a larger amount of scatter than their nominal 1-σ uncertainties allow.

We augmented `tempo` to include a parameter that acts analogously to EFAC[35], which we call "DMEFAC"[36]. This parameter was tuned iteratively until the reduced

[35] Compared to Equation 4.6, $\sigma_{\mathrm{DM},i} \rightarrow \mathrm{DMEFAC} \times \sigma_{\mathrm{DM},i}$.
[36] Thanks to P. Demorest.

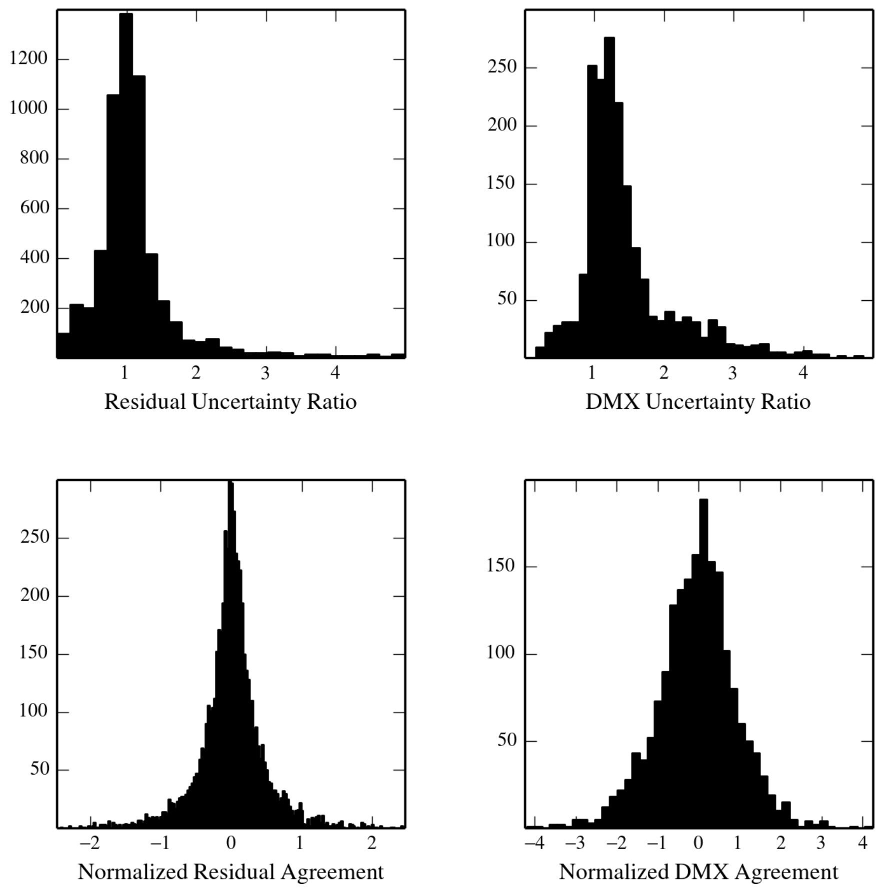

Fig. 4.7 – Histograms of the information contained in the comparison plots between Figures 4.46 & 4.118. The bottom row contains the normalized differences between the residuals and DMX values as measured here and in A15; the top row contains the ratio of the uncertainties on these quantities (wideband/A15). See text for details.

χ^2 value was within 0.01 of 1.0, as can be seen in the table. The DMEFACs are listed; 22 out of 37 are < 2, another 12 are < 4, and the remaining three are $\sim$6, 6, and 14. The pulsars with the most extreme DMEFACs, however, are all red noise pulsars. Therefore, because DMEFAC was not adjusted based on the whitened residuals, these cannot be seen as inaccuracies in our wideband DM uncertainty estimates. 9 non-red noise pulsars have DMEFACs between 2.0 and 4.0. The DMEFAC parameters are *not* incorporated into the timing summary figures (though the properly modeled noise parameters are).

The last two pairs of columns in Table 4.5 show the weighted root-mean-square (WRMS) residual level, which is a typical figure of merit for assessing an MSP's quality. For our "whitened" WRMS residuals, we subtracted the maximum likelihood realization of the red noise and then calculated the WRMS as

$$\mathrm{WRMS} = \sum_{i=1}^{\#\,\mathrm{resids}} \left(\frac{\delta t_i - \overline{\delta t}}{\sigma_i}\right)^2 \Bigg/ \sum_{i=1}^{\#\,\mathrm{resids}} \frac{1}{\sigma_i^2}, \tag{4.9}$$

where i indexes the residual δt, which has noise-modeled uncertainty σ, and the average residual $\overline{\delta t}$ is itself a weighted average,

$$\overline{\delta t} = \sum_{i=1}^{\#\,\mathrm{resids}} \frac{\delta t_i}{\sigma_i^2} \Bigg/ \sum_{i=1}^{\#\,\mathrm{resids}} \frac{1}{\sigma_i^2}. \tag{4.10}$$

Many pulsars show very comparable timing results (e.g., J0023+0923, J1614−2230, J2145−0750, J2317+1439), differing by no more than a few to ∼10 percent. For these pulsars, we can be fairly confident in our measurements. In other cases, it is difficult to say whether or not the improvement (or worsening) of the timing residual is due to our portrait modeling, wideband measurements, or if it is an artifact of the noise modeling. The very large differences, particularly the large improvements (e.g., J1853+1303) must have origins in the noise parameters; there is no reason for the profile modeling or the effective TOA averaging to influence the results at this level. For this reason, we need to take these results, and the timing parameter results that follow, with caution.

The only pulsar whose (white) WRMS residual got significantly worse (by almost a factor of two) is J1923+2515, whose model was discussed in §4.2.1. The only pulsar whose non-whitened residual grew significantly was J0613−0200, which is reflected in the significantly steeper red noise index that we measured. The next two pulsars with worse WRMS values are J1455−3330 and J1741+1351, both at the level of ∼50%.

Of the four pulsars where we modeled the scattering, three of them showed significant improvement in the timing residuals (whitened or not), and the fourth (J1600−3053, which has the smallest scattering timescale of the four) remained at the same WRMS level. Approximately 70% of the pulsars showed some or no improvement in the WRMS timing residual.

The last thing to consider is how the fitted timing model parameters differed. Figures 4.119–4.155 show the difference in each fitted astrometric, spin, and binary parameter, normalized by the uncertainty from A15, per pulsar (in faux "σ" units). The size of the "error bar" corresponds to the ratio of the uncertainties, where > 1 means that our uncertainties are larger. These two quantities are given in the plot. Each point is colored based on a 3-"σ" significance threshold: black points exceed this threshold in both analyses, blue points fall below this threshold in both cases[37], red points are where a parameter is over the threshold only in the wideband analysis, and green points are where a parameter is over the threshold only in the A15 analysis.

Almost across the board, the parameters do not differ by more than 2–3 "σ". In only a few cases there are there larger deviations: J1713+0747, J1744−1134 and J2017+0603. The only parameters that appear red or green (i.e., that switch significance) are timing model parameters that are marginally significant in the first place, or require long timing baselines, or are highly non-linear: binary derivatives, long-term binary parameters, parallax, proper motion, Shapiro delay. One of the apparently "new" parallax measurements is probably artificial since we don't have a reasonably long baseline (J1832−0836 has only a 0.6 yr baseline). Similarly, J0931−1902 appears to have "lost" the significance of its spin down parameter (F1), but this is probably

[37]In particular, parallax was kept in the A15 timing models, regardless of significance. The same is true for proper motion, except that it wasn't fit for either J0931−1902 or J1832−0836, both of which have <1 yr of data.

Table 4.5. Timing Results & Comparison

				Full WRMS [μs]		Whitened WRMS [μs]	
PSR	χ^2	$\mathrm{N_{DoF}}$	DMEFAC	Wideband	A15	Wideband	A15
J0023+0923	414.04	413	1.40	0.354	0.320	-	-
J0030+0451	227.99	227	1.68	0.715	0.723	0.078	0.212
J0340+4130	82.87	83	1.21	0.301	0.385	-	-
J0613−0200	384.61	381	1.37	1.448	0.592	0.229	0.165
J0645+5158	96.80	96	1.26	0.043	0.052	-	-
J0931−1902	29.01	29	1.23	0.281	0.381	-	-
J1012+5307	679.17	674	1.11	0.709	1.197	0.254	0.355
J1024−0719	167.35	167	1.15	0.157	0.280	-	-
J1455−3330	325.66	323	1.21	1.052	0.694	-	-
J1600−3053	307.59	306	2.00	0.203	0.197	-	-
J1614−2230	281.89	281	1.07	0.192	0.189	-	-
J1640+2224	290.78	291	3.50	0.077	0.158	-	-
J1643−1224	346.08	343	2.35	1.985	2.058	0.095	0.331
J1713+0747	926.50	923	3.70	0.147	0.116	-	-
J1738+0333	198.38	198	1.43	0.257	0.308	-	-
J1741+1351	107.72	107	1.70	0.159	0.103	-	-
J1744−1134	455.79	455	3.00	0.264	0.334	-	-
J1747−4036	67.47	67	2.53	0.353	0.531	-	-
J1832−0836	29.94	30	1.23	0.084	0.121	-	-
J1853+1303	105.15	105	1.25	0.040	0.235	-	-
B1855+09	385.06	383	2.33	1.267	1.338	0.414	0.505
J1903+0327	90.30	90	2.20	1.401	1.949	0.041	0.327
J1909−3744	514.37	511	1.82	0.063	0.081	-	-
J1910+1256	182.41	183	1.90	1.347	1.449	0.271	0.587
J1918−0642	528.21	526	1.23	0.318	0.340	-	-
J1923+2515	50.45	50	3.45	0.496	0.266	-	-
B1937+21	622.27	624	14.40	1.536	1.550	0.128	0.104
J1944+0907	117.03	116	1.97	1.148	2.440	0.413	0.332
J1949+3106	56.22	56	2.50	0.726	0.646	-	-
B1953+29	102.26	102	6.40	5.971	4.149	0.443	0.531
J2010−1323	262.00	261	1.28	0.297	0.312	-	-
J2017+0603	83.18	83	1.41	0.057	0.073	-	-
J2043+1711	114.30	114	1.14	0.132	0.108	-	-
J2145−0750	285.13	283	2.52	0.364	0.370	-	-
J2214+3000	135.58	135	2.50	0.258	0.314	-	-
J2302+4442	85.23	85	1.45	0.769	0.708	-	-
J2317+1439	430.41	427	1.73	0.259	0.267	-	-

Note. — The reduced χ^2 values are all very near 1.0 due to the noise modeling in combination with the listed DMEFAC parameter. The only large DMEFACs arise in pulsars with red noise, which doesn't reflect the proper DMEFAC. The columns labeled A15 are from Table 3 of Arzoumanian et al. (2015a).

a combination of the noise parameters over-inflating the uncertainties and its short timing baseline (also 0.6 yr). The other "new" parallax measurements (J1640+2224, J1741+1351, J1747−4036, with time spans of 8.9, 4.2, and 1.7 yr, respectively) may be credible, but again, it will depend on how robust the noise parameters are.

4.3 Conclusion

In this chapter, we have further developed our wideband timing methods and applied them broadly to a sample of 37 relatively bright, well-timed MSPs monitored by the North American Nanohertz Observatory for Gravitational Waves (NANOGrav). Specifically, we performed an analogous timing analysis on the same underlying data presented in Arzoumanian et al. (2015a) (A15). All TOAs and timing models used in these analyses from A15 are available on the internet[38]. The analogous timing results from our wideband analyses, as well as Gaussian models and numerous plots from this chapter, are also available[39].

To summarize the main results and findings from our analyses:

1. We made Gaussian component models for all pulsars to model profile evolution with frequency, choosing between several combinations of evolutionary models.
2. Wideband timing measurements were made for the entire data set. Constant DM offsets reflecting small modeling errors were measured and subtracted.
3. A comprehensive timing analysis was performed, including a full Bayesian noise analysis. We augmented `tempo` to include our wideband DM measurements as data for the DM(t) model.
4. We recover virtually identical red noise parameters in all but one or two pulsars, which is a significant test to pass if these measurements are to be used in PTA experiments.

[38] `http://data.nanograv.org/`
[39] `http://www.astro.virginia.edu/~ttp4tx/nanograv/`

5. Our in-band DM measurements may help remove chromatic red noise trends seen in the residuals of some pulsars and may indicate the importance of characterizing a frequency dependent DM.
6. In most cases, we see the same or improved levels of WRMS timing residual, although these quantities — as well as the fitted timing parameters — are subject to further scrutiny based on a more judicious selection of noise parameters.
7. The differences in the fitted astrometric, spin, binary, and DMX parameters are typically negligible, with the caveat mentioned above.

With respect to the first item, if one only cares about characterizing the data, Gaussian models will perform reasonably well for pulsars with uncomplicated profile evolution that is easily represented by Gaussian functions, or in sufficiently bright pulsars such that one can disentangle dispersive delays and still model profile evolution. The Gaussian components can contain interesting information regarding intrinsic properties of the magnetosphere, but the focus will remain on simply characterizing the data until truly broadband studies and receivers are the norm since much of the profile evolution still occurs in the gaps between frequency bands.

Other modeling approaches (e.g., using Principle Component Analysis, wavelets, interpolation, etc.) should also be explored. Our immediate goal is to simply extend the family of evolutionary models used; in particular, to use a "Thorsett" model of a power-law + constant (Thorsett 1991). Automated modeling algorithms are a more distant goal but one possible approach is to use a technique called Reversible-Jump Markov Chain Monte Carlo, which allows sampling between posteriors of different dimensions. In this way, one could intelligently choose between different families of Gaussian (or other) models with different numbers of components, etc.

An upcoming version of the `PAL2` noise modeling software will include DMDATA functionality. Along with the TOAs, the calculation of the likelihood will make use of the in-band measured DMs, and model at least basic white noise parameters like

DMEFAC and potentially "DMEQUAD". A new "DMJUMP" parameter in `tempo` would be a natural way to correct unmodeled profile evolution for the time being. Eventually, we will integrate an algorithm to simultaneously measure the TOA, DM, and scattering timescale, instead of having the latter parameter be a fixed component of the model.

Some of our findings here can be validated with simulations of profile evolution compounded with the other effects discussed: DISS (Cordes in prep.), DM(ν) (Cordes et al. 2015), a variable scattering timescale, etc. In particular, we wish to know whether or not we have successfully removed chromatic dependencies in the residuals of some high DM pulsars, and perhaps lessened the significance of shallow red noise in others. This latter question can be answered simply by calculating the Bayes factors for our noise models.

It is apparent that we have additional developments to make in order to fully supplant current techniques in the era of truly broadband receiver systems, but we have demonstrated that our wideband methodology performs at least as well in many cases, if not better. Although in principle the wideband method should never do any worse than current methods, it is obvious that the next set of developments will continue to naturally cater to the future of pulsar astronomy. This is in contrast to what would have to be only additionally *ad hoc* methods if channelized TOAs and frequency-independent profile templates continue to be used. For example, the number of FD parameters would grow significantly with increasing bandwidths, whereas in the worst case our DMJUMP parameters are limited by the number of receivers used. The simultaneous use of the full timing band from, e.g., 0.6–3 GHz would improve the DM determination, which would in turn allow for more constraining DM(t) models.

Even in the absence of improved models or noise analyses, it will be an interesting exercise to run our timing residuals through the same GW detection pipeline that will soon analyze the A15 results. If nothing else, the problem will be much easier due to the data set being a factor of $\sim$30 smaller, providing results hundreds of times faster. The long term goal, however, is to optimize both the strategy and sensitivity of PTA experiments.

4.4 Appendix: Per-Pulsar Plots

Here, we append plots relevant to the discussions in this chapter, but include them for all pulsars. The color legend indicating the receiver from which a measurement was made is originally in Table 4.2: red = 327 MHz, orange = 430 MHz, green = Rcvr_800, blue = Rcvr1_2, blue (or dark blue, in frequency residual plots) = L-wide, purple = S-wide. To recapitulate the plots and their primary contents:

- Figures 4.8–4.44 contain the post-fit frequency residuals. Constant offsets have been removed so that all bands lie on the zero line. Each figure shows what the average timing deviation is as a function of frequency from assuming a constant profile shape in each separate band. Note that some profile evolution gets absorbed by the joint fit (across bands) for a DM with DMX. When fit together with FD parameters, these systematics flatten out. The size of the "residual DM" as measured by a ν^{-2} law fit to the residuals *in each band* is a fair metric by which we can asses the size of our DM offsets in Table 4.4 and Figure 4.5. The number of FD parameters used in A15 to characterize these trends are also given in the table. Note that the excessive scatter/outliers in some pulsars is a result of either persistent, unzapped RFI, or a pulsar "ghost" artifact signal from imperfect sampling of the raw data signal. Also note that the two lowest frequency bands are compressed significantly, so their trends may be suppressed graphically here.

- Every-other figure starting with Figure 4.45 and ending with Figure 4.117 contains a timing summary plot. These are analogous to what is published in A15. The topmost panel in each figure shows the wideband timing residuals for all nine years 2005–2014, indicating when the switch in backends occurred.

The bottommost panel shows the wideband DMX measurements. If the pulsar has detected red noise, the second topmost panel shows the whitened residuals, which are explained in the text. The second bottommost panel shows the mean-subtracted wideband ("in-band") DM measurements *on the same scale* as the DMX DMs. Because the wideband DMs from ASP & GASP are usually completely uninformative, we have not plotted their error bars to reduce clutter. Also note that the wideband DM uncertainties do not incorporate DMEFAC as discussed in §4.2.3. All other uncertainties are plotted with the noise parameters included.

- Every-other figure starting with Figure 4.46 and ending with Figure 4.118 contains a timing comparison plot. The dotted lines in both subplots represent equality. In the left subplot, "Standard Residual" means the averaged residuals from fitting the channelized TOAs in A15; i.e., these are the data from the middle panel of the analogous plot in A15. For pulsars with red noise, these are the whitened residuals. Similarly, in the right subplot, "Standard DMX" contains the corresponding DMX values from A15. Both axes on each subplot have the same scale. All uncertainties are plotted with the noise parameters included.

- Figures 4.119–4.155 contain plots that compare the timing model parameter values between our analysis and A15. As mentioned in the text, the plotted points are the difference between the values, normalized by the corresponding A15 uncertainty (and so are in faux units of "σ"). The uncertainties on the points are the ratio of our parameter uncertainties to those from A15. These two numbers are displayed in the plot. The colors are related to a significance threshold of 3-"σ": black points exceed this threshold in both analyses, blue points fall below this threshold in both cases[40], red points are where a parameter is over the threshold only in the wideband analysis, and green points are where a parameter is over the threshold only in the A15 analysis. The astrometric parameters are: LAMBDA = ecliptic longitude, BETA = ecliptic latitude, PMLAMBDA = proper motion in the LAMBDA direction, PMBETA = proper motion in the BETA direction, PX = parallax. The spin parameters are: F0 = spin frequency, F1 = spin-down frequency. The binary parameters for a generic orbit[41] are: A1 = size of the projected semi-major axis, PB = orbital period, OM = longitude of periastron, T0 = epoch of periastron passage, E = orbital eccentricity. The binary parameters for a low-eccentricity orbit[42] become: the Laplace-Lagrange parameters EPS1 = $\mathrm{E}\sin(\mathrm{OM})$ and EPS2 = $\mathrm{E}\cos(\mathrm{OM})$, TASC = epoch of the

[40] See footnote 37.
[41] e.g., The "BT" or "DD" models in `tempo`.
[42] e.g., The "ELL1" model in `tempo`.

ascending node = T0−(OM × PB/2π). The binary evolution parameters are: XDOT = derivative of A1, PBDOT = derivative of PB, OMDOT = derivative of OM, EPS1DOT = derivative of EPS1, EPS2DOT = derivative of EPS2. The Shapiro delay parameters are: SINI = sin of the inclination of the orbit, M2 = mass of the companion object.

- Every-other figure starting with Figure 4.156 and ending with Figure 4.228 contains concatenated data portraits from each band. In all cases except for the four mentioned in §4.2.1, the phase-frequency portrait in the left panel consists of quasi-coherently averaged data from several high S/N epochs. The receiver band edges are provided in Table 4.2, but can also be seen in the panels on the right. The other frequency gaps in the left panel are channels that have been removed because they contained persistent RFI. The data were aligned based on the fitted Gaussian model, which determines DMs and relative phase offsets to align the data. The top-left sub-panel contains the average profile, summed over the entire range of frequencies. The dotted line is the zero-baseline level and the horizontal line below this is a scale of length 1 ms, with the innermost thicker box being a scale of length 100 μs. The left-most sub-panel plots the phase-average of each channel, where each channel *has been normalized by the peak value of that profile.* Therefore, the "spectral index" γ at the top of this sub-panel reflects something like how the overall duty-cycle/pulsed area changes with frequency, *not the flux density*. The panels on the right contain the same aligned data as in the left panel, but they are rescaled. The frequency bands are separated and the channels zapped of RFI are not plotted. The size of vertical space that a band occupies is proportional to the size of the dispersive delay across the individual band relative to the total dispersive delay between the bottom of all bands to the top of all bands. This scaling compensates for the linear spacing in the left panel where the low frequency bands that carry a lot of dispersive delay and profile evolution appear narrow. The top-right sub-panel contains average profiles color-coded by the band from which they were averaged. The color bar applies to all phase-frequency panels.

- Every-other figure starting with Figure 4.157 and ending with Figure 4.229 contains the fitted Gaussian component model and residuals. The left panel shows the model constructed to span the same range as the left panel of the corresponding portrait plot; the color bar is also the same. The residuals from subtracting these two panels are in the right panel, which have their own color bar. In a number of pulsars, you can see residual RFI and, in some cases, data acquisition artifacts. Nominally, the wideband fitting method should be relatively immune to these imperfections, whereas they can significantly bias certain channelized TOAs.

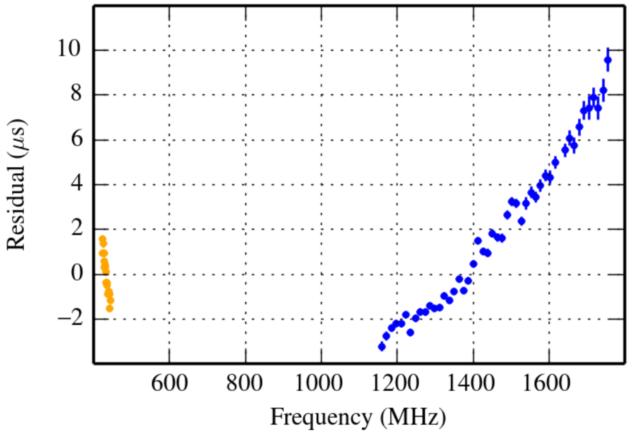

Fig. 4.8 – J0023+0923

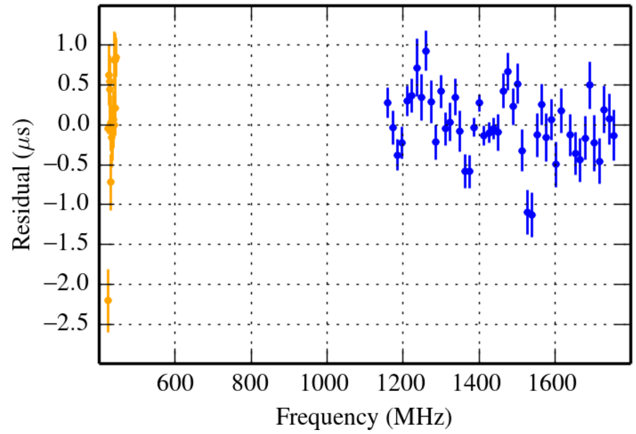

Fig. 4.9 – J0030+0451

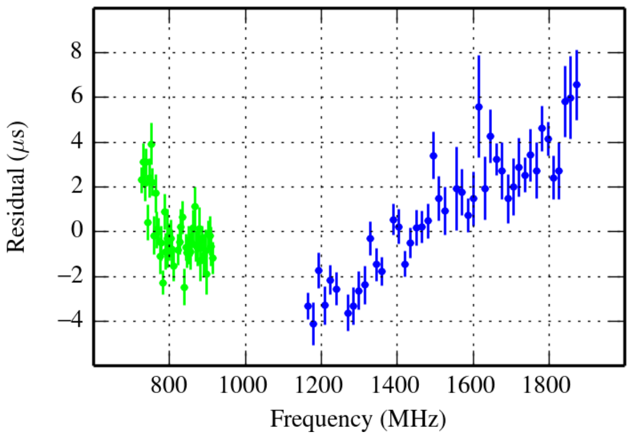

Fig. 4.10 – J0340+4130

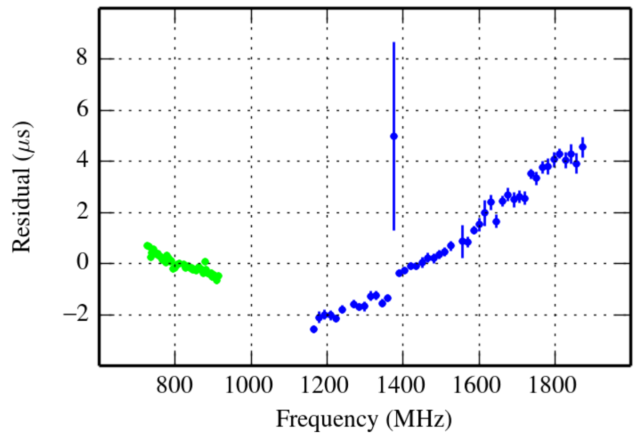

Fig. 4.11 – J0613−0200

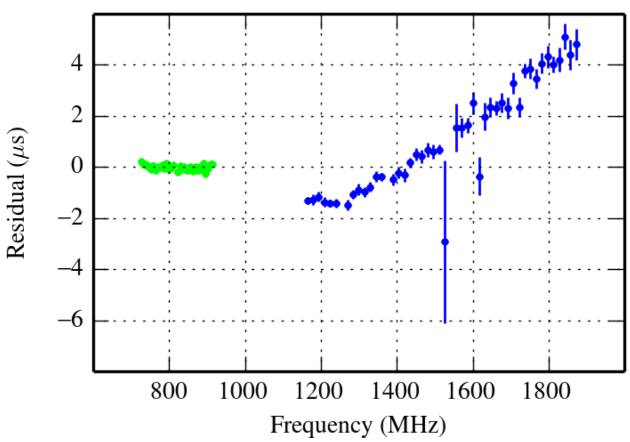

Fig. 4.12 – J0645+5158

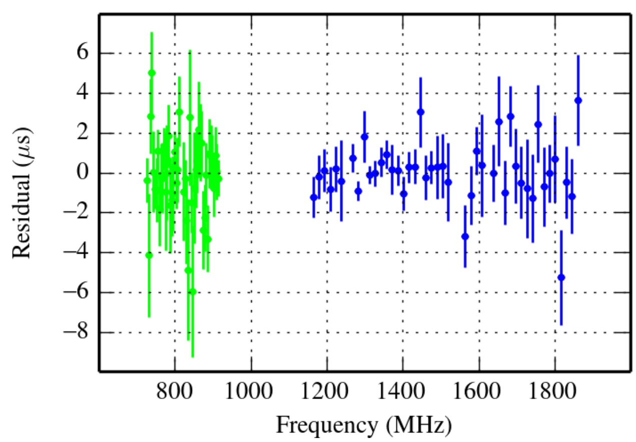

Fig. 4.13 – J0931−1902

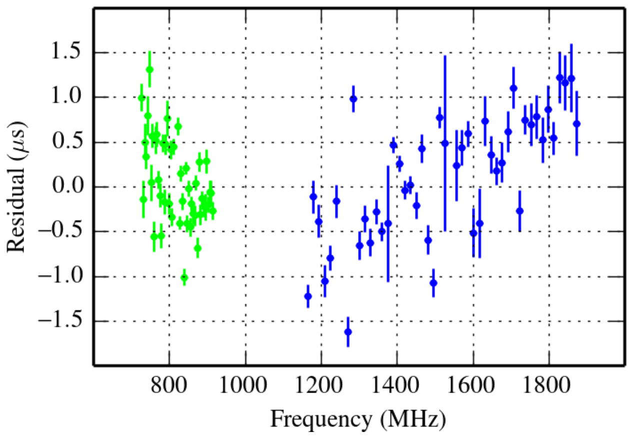

Fig. 4.14 – J1012+5307

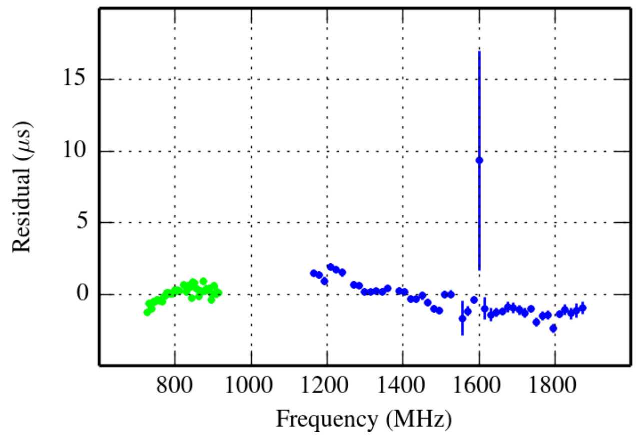

Fig. 4.15 – J1024−0719

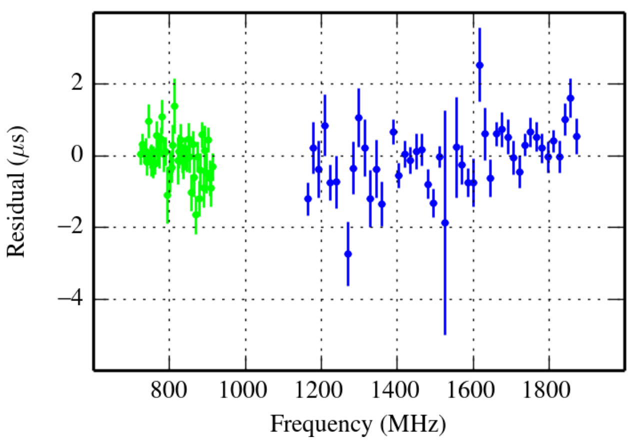

Fig. 4.16 – J1455−3330

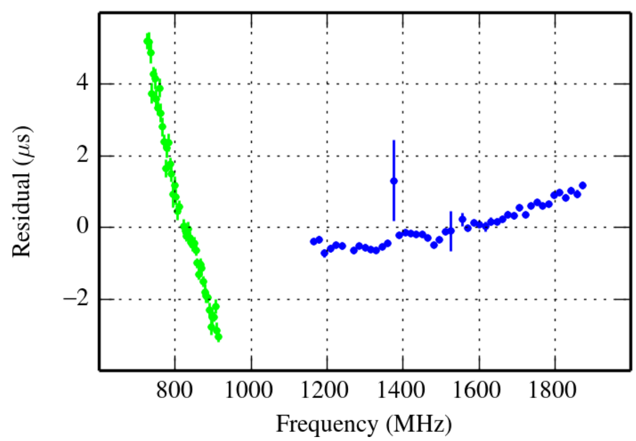

Fig. 4.17 – J1600−3053

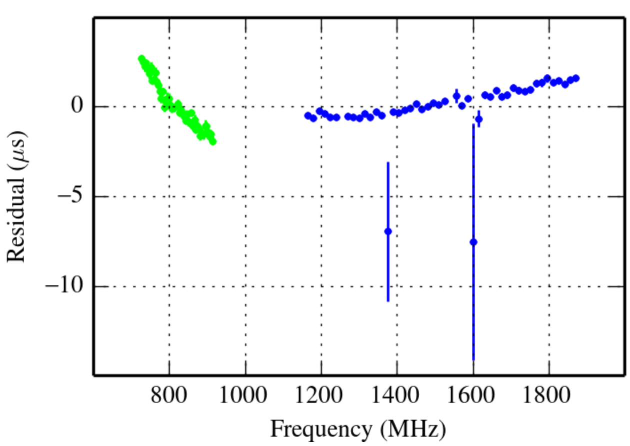

Fig. 4.18 – J1614−2230

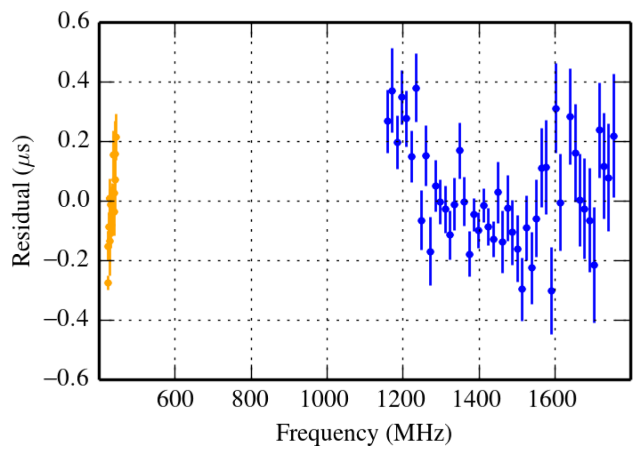

Fig. 4.19 – J1640+2224

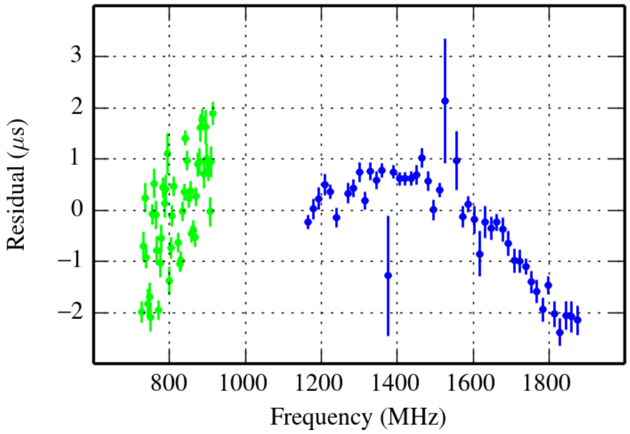

Fig. 4.20 – J1643−1224

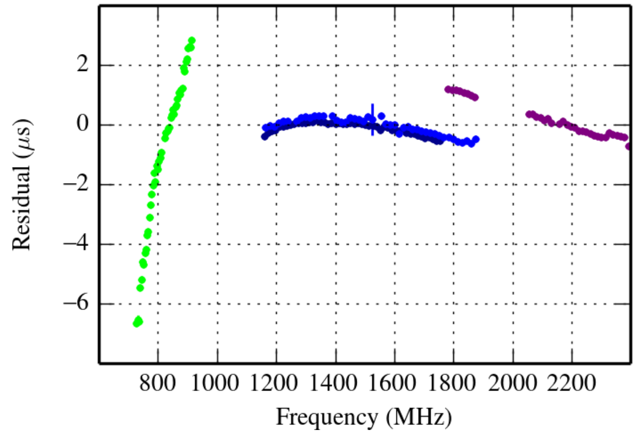

Fig. 4.21 – J1713+0747

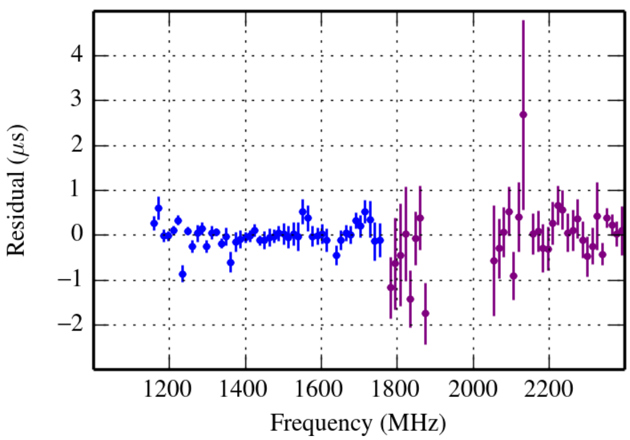

Fig. 4.22 – J1738+0333

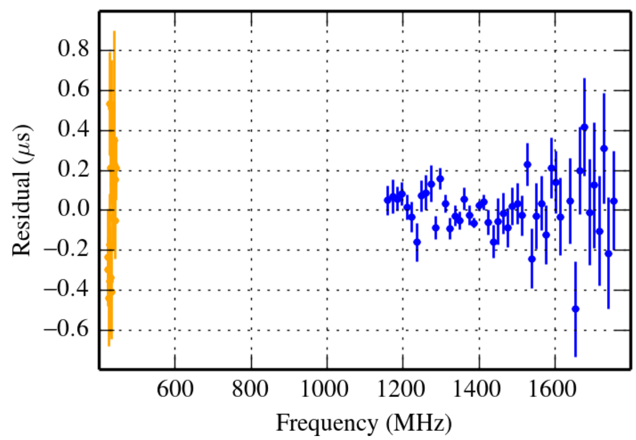

Fig. 4.23 – J1741+1351

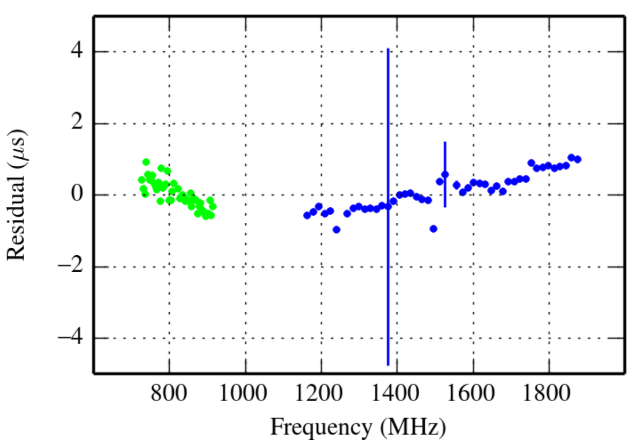

Fig. 4.24 – J1744−1134

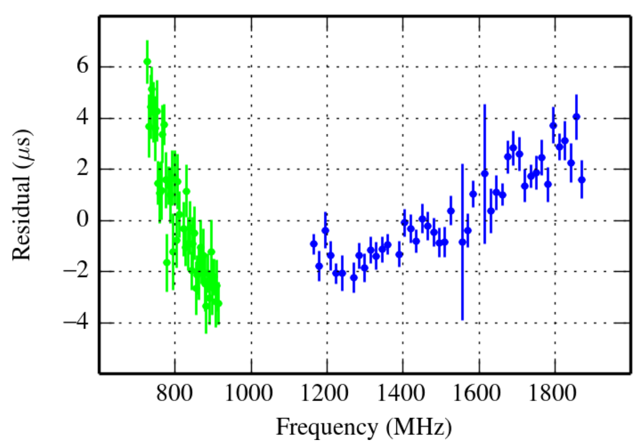

Fig. 4.25 – J1747−4036

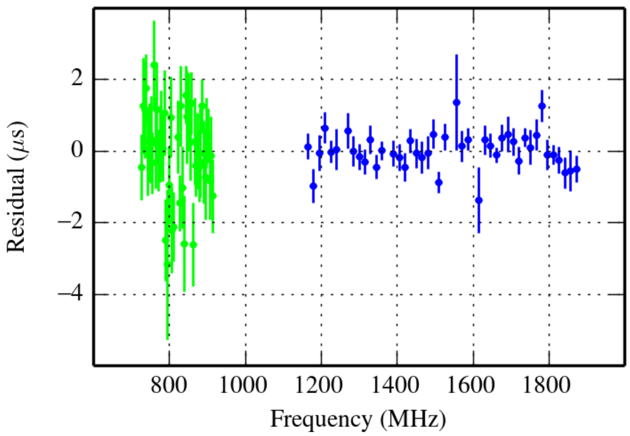

Fig. 4.26 – J1832−0836

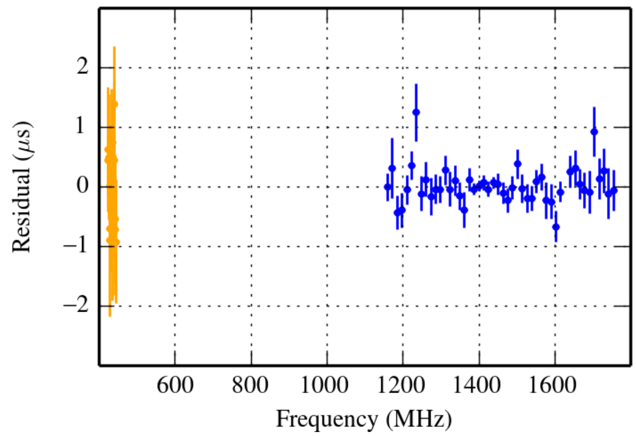

Fig. 4.27 – J1853+1303

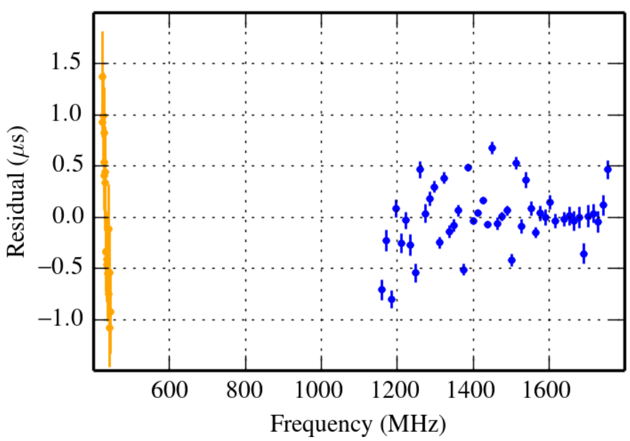

Fig. 4.28 – B1855+09

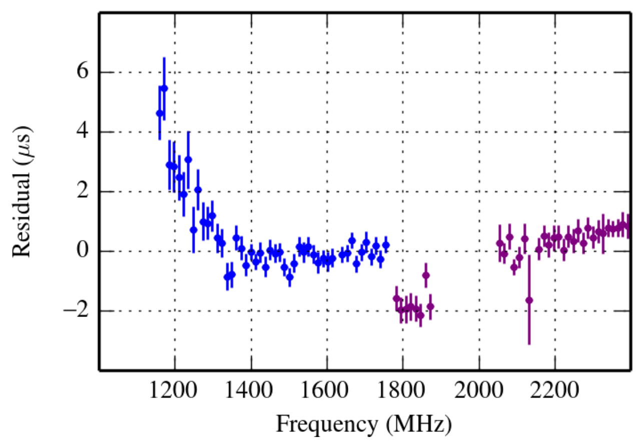

Fig. 4.29 – J1903+0327

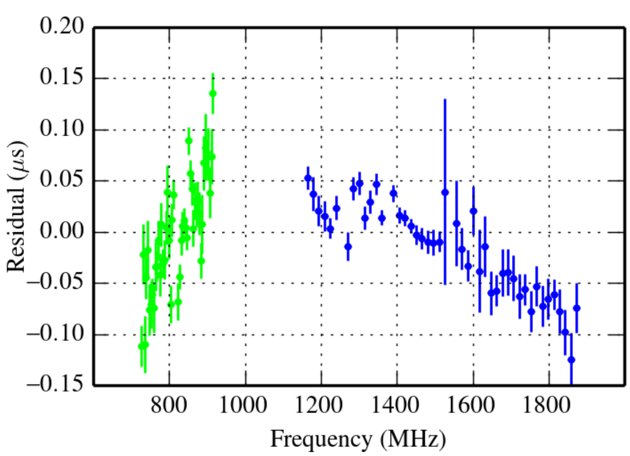

Fig. 4.30 – J1909−3744

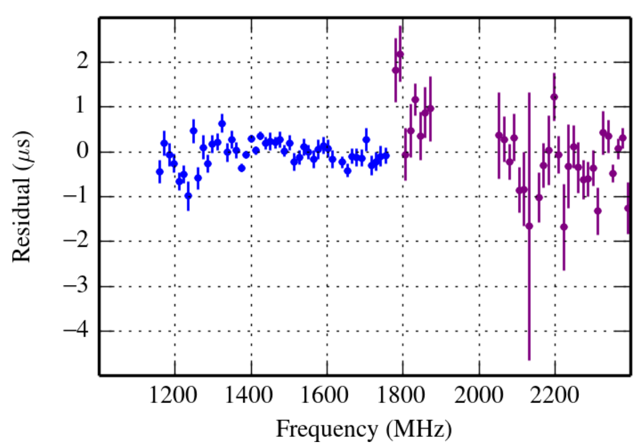

Fig. 4.31 – J1910+1256

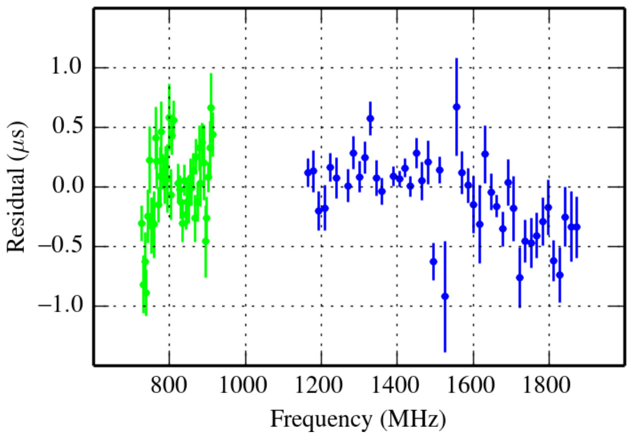

Fig. 4.32 – J1918−0642

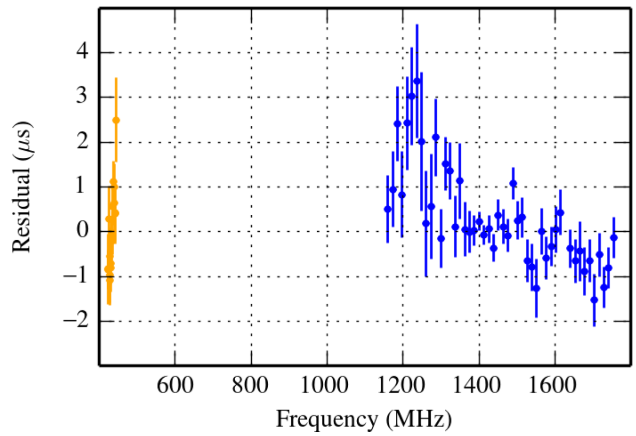

Fig. 4.33 – J1923+2515

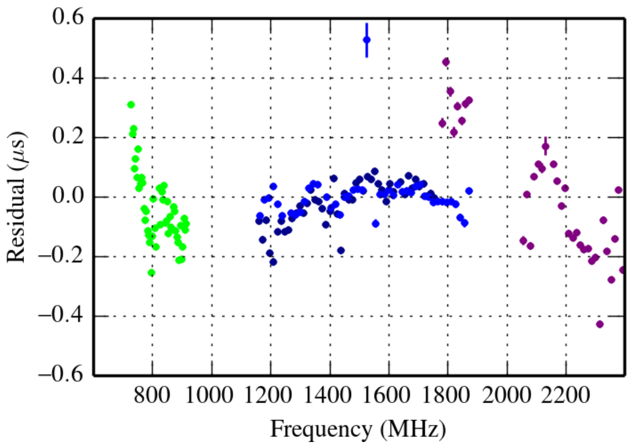

Fig. 4.34 – B1937+21

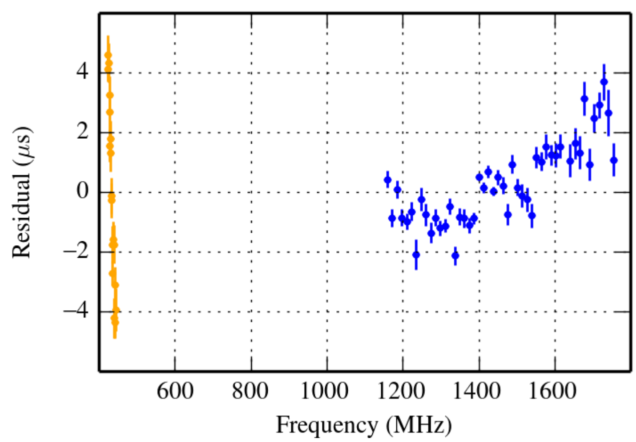

Fig. 4.35 – J1944+0907

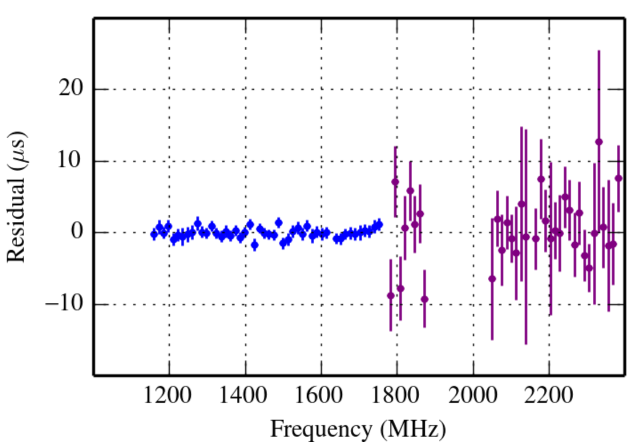

Fig. 4.36 – J1949+3106

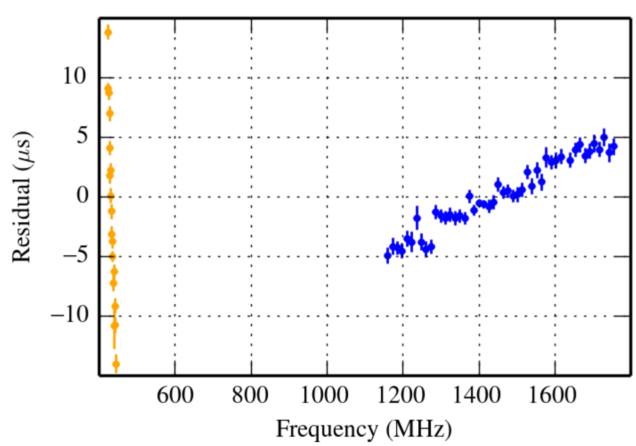

Fig. 4.37 – B1953+29

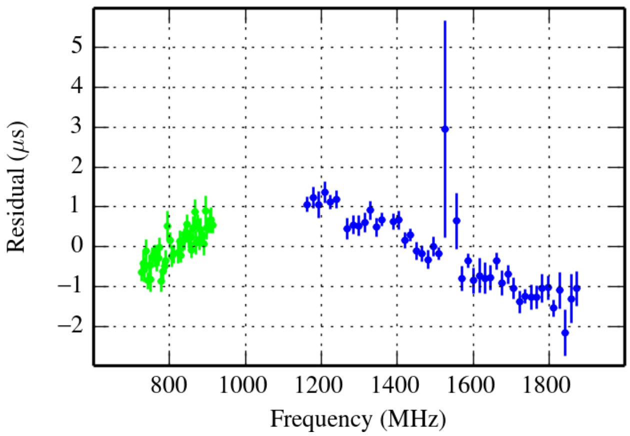

Fig. 4.38 – J2010−1323

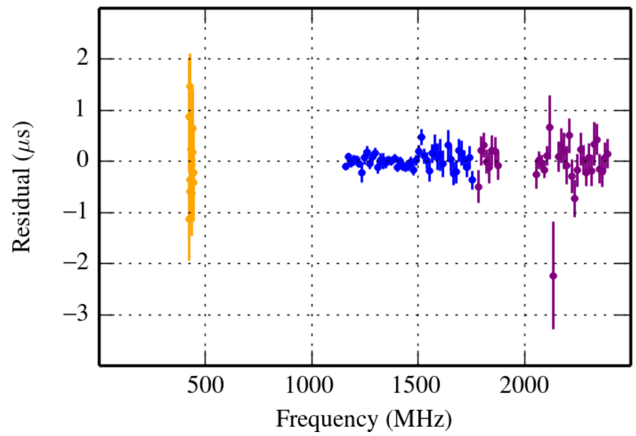

Fig. 4.39 – J2017+0603

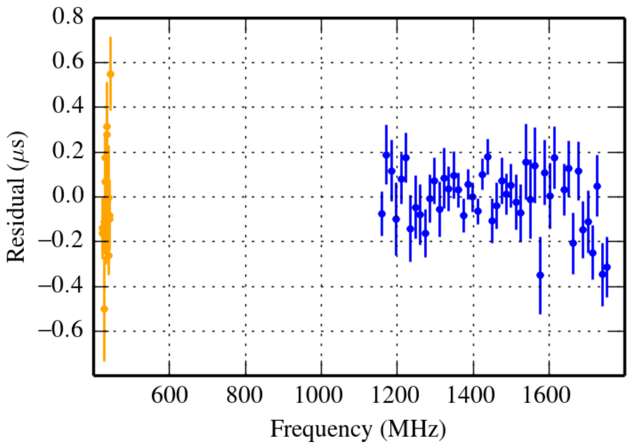

Fig. 4.40 – J2043+1711

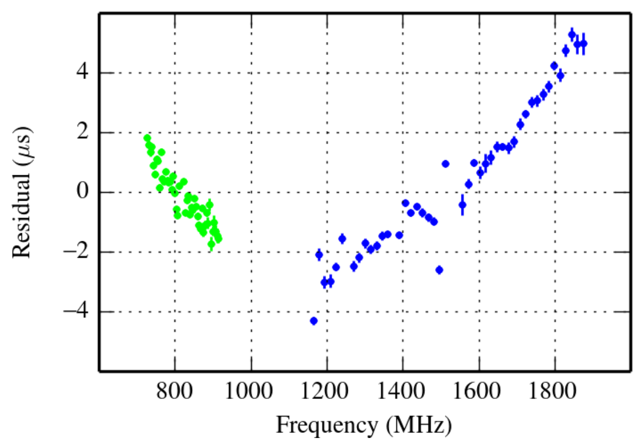

Fig. 4.41 – J2145−0750

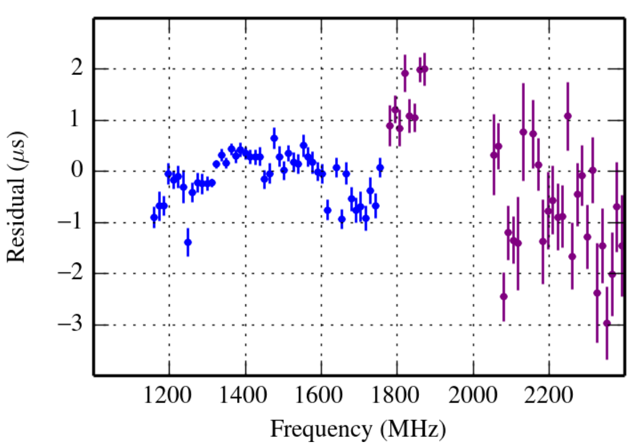

Fig. 4.42 – J2214+3000

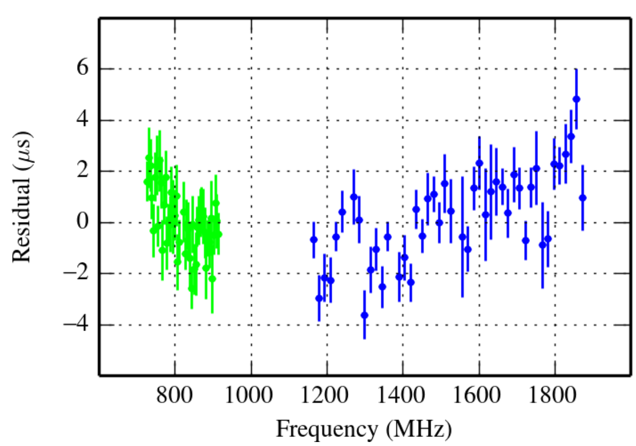

Fig. 4.43 – J2302+4442

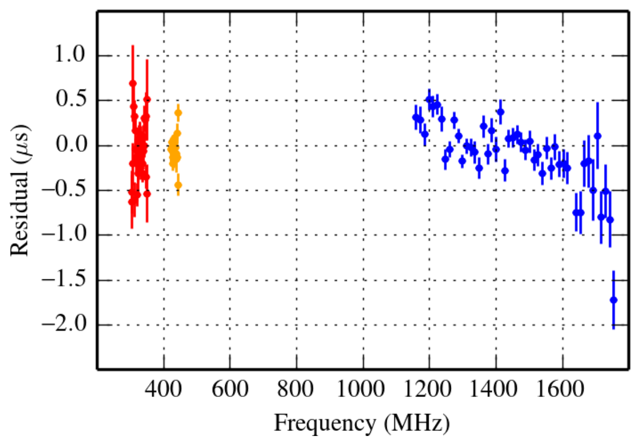

Fig. 4.44 – J2317+1439

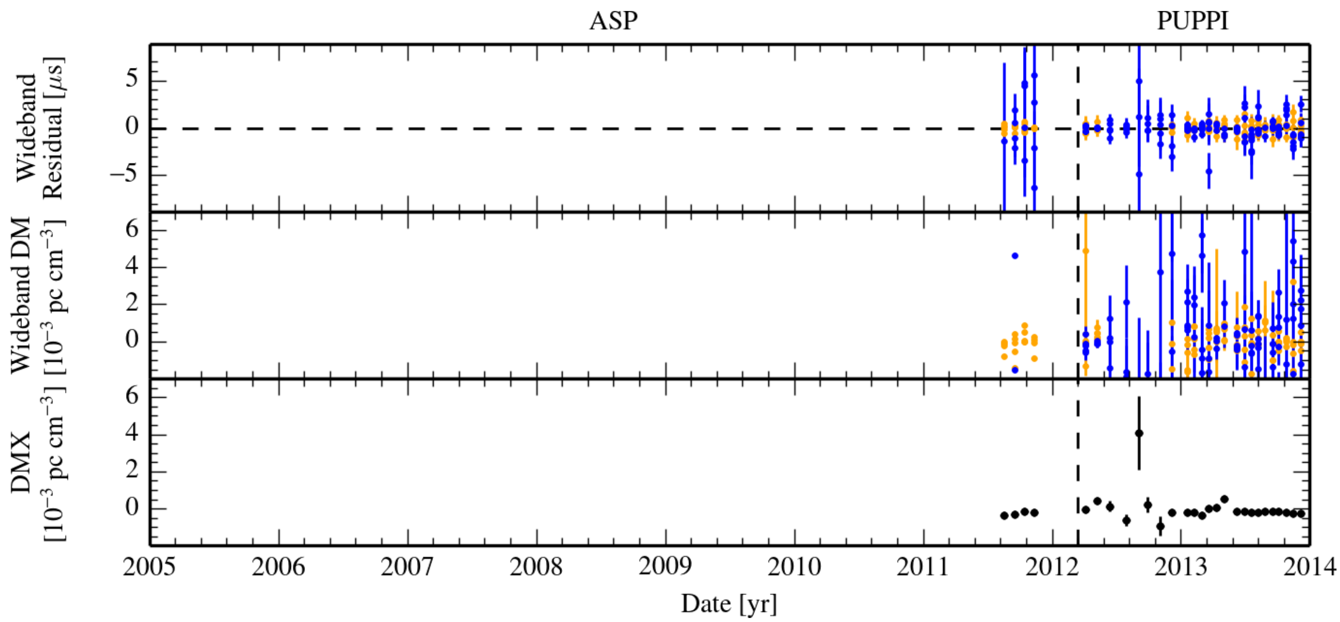

Fig. 4.45 – Timing summary panels for J0023+0923.

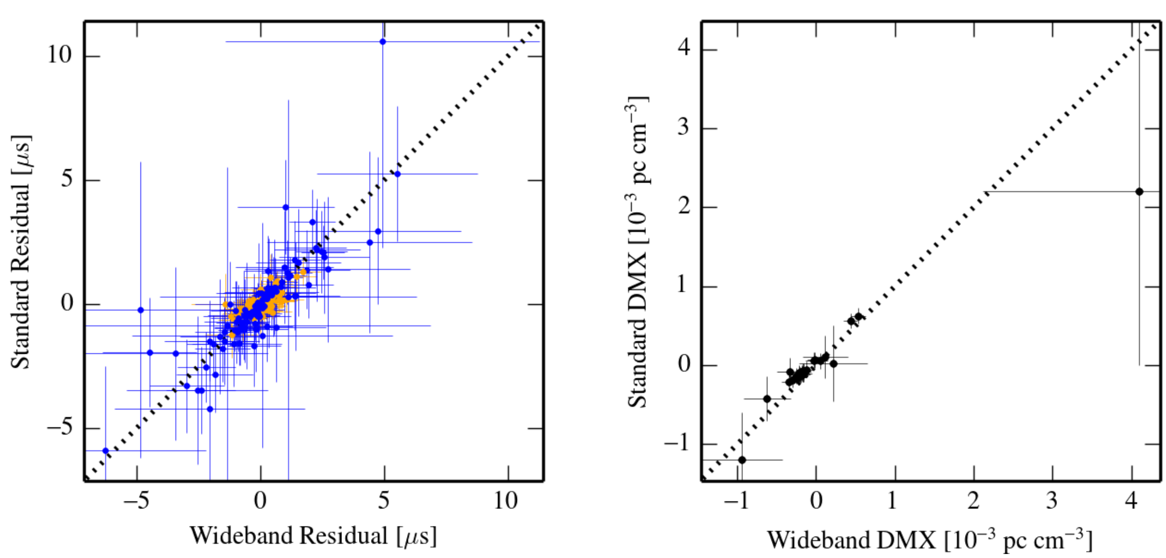

Fig. 4.46 – Comparison subplots for J0023+0923.

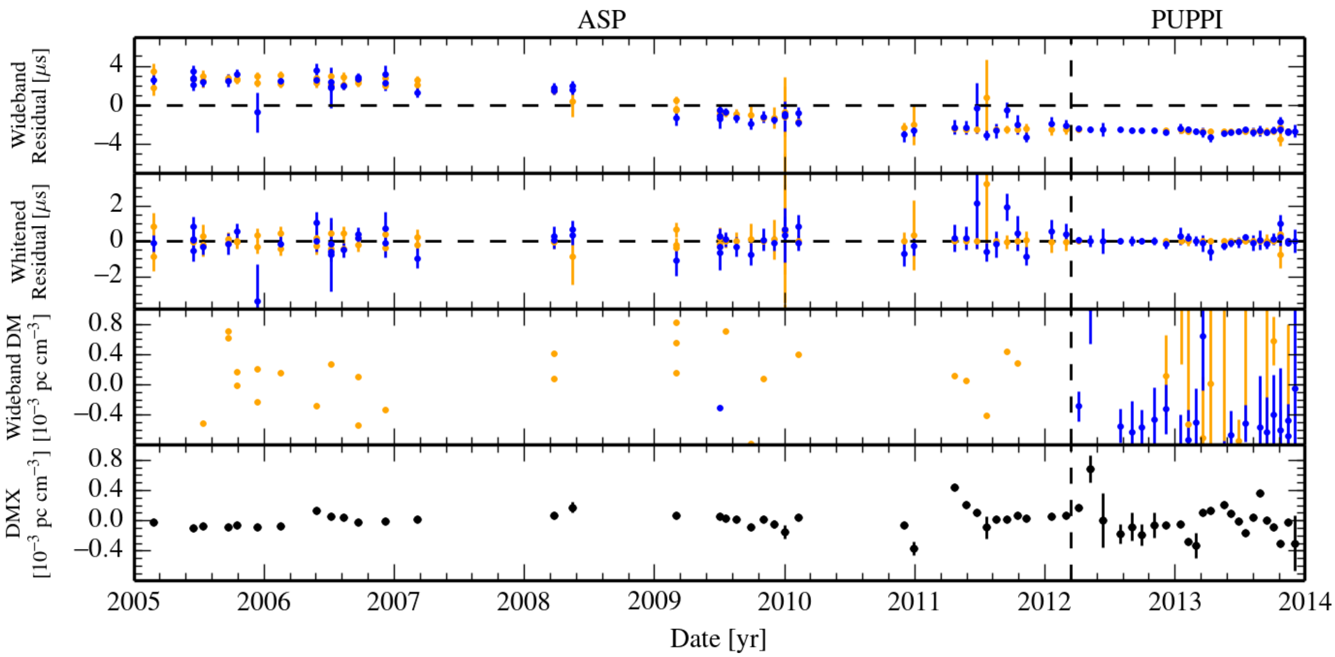

Fig. 4.47 – Timing summary panels for J0030+0451.

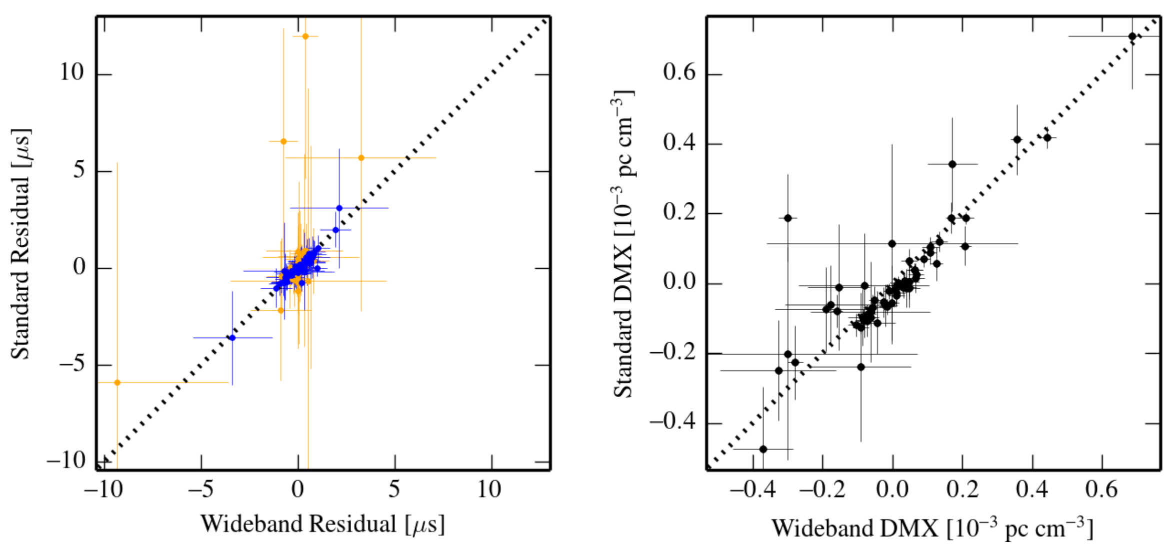

Fig. 4.48 – Comparison subplots for J0030+0451.

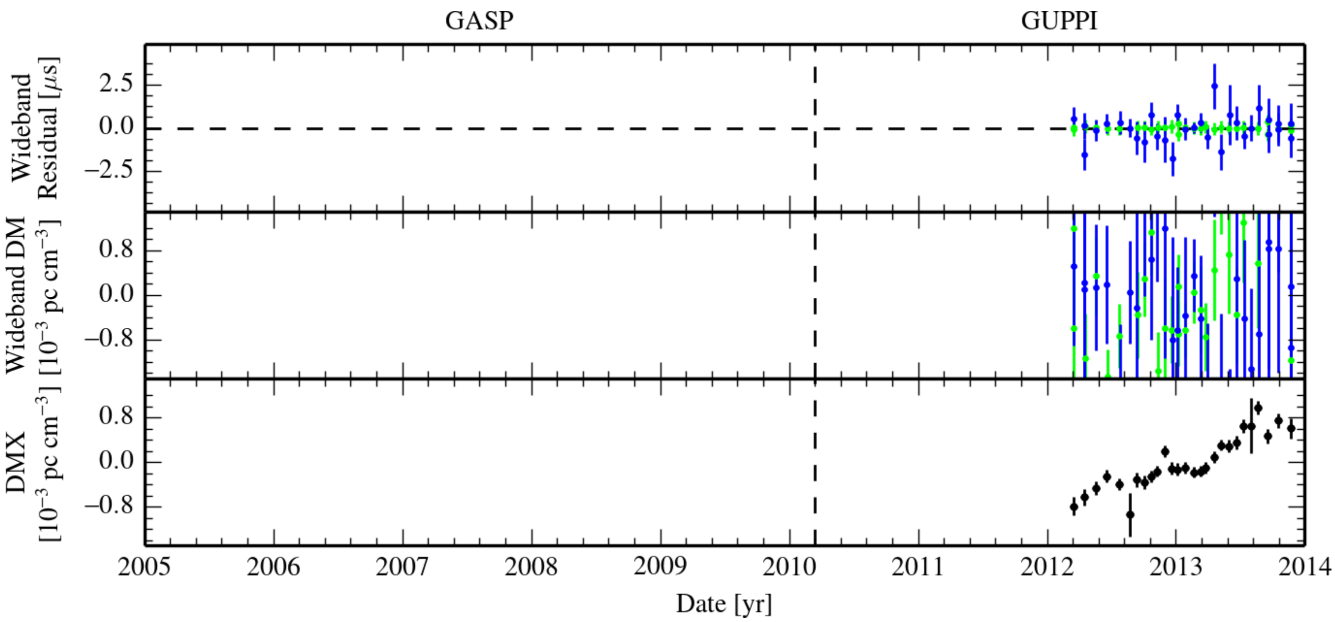

Fig. 4.49 – Timing summary panels for J0340+4130.

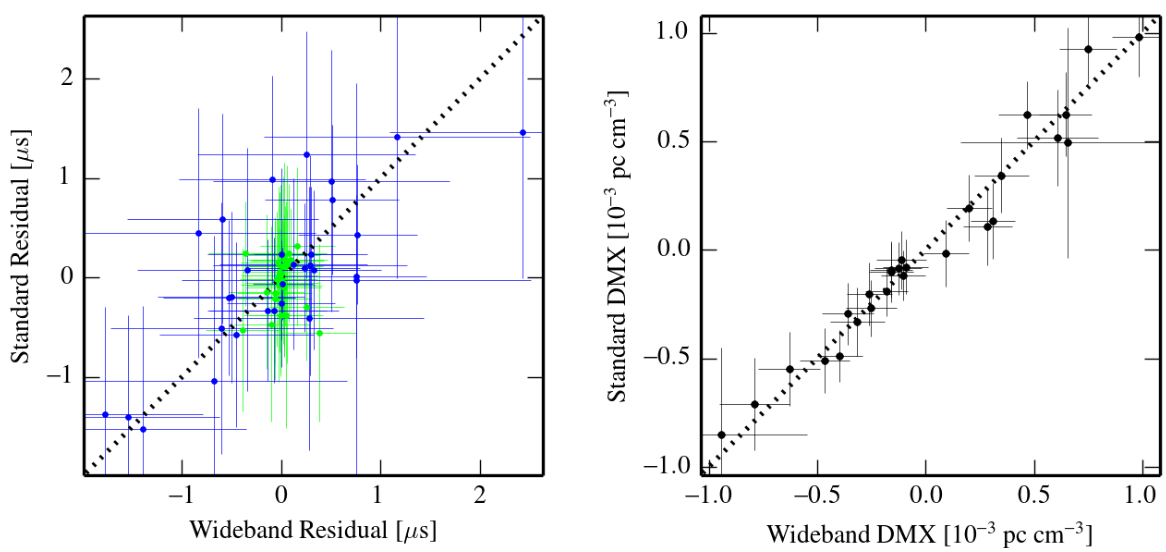

Fig. 4.50 – Comparison subplots for J0340+4130.

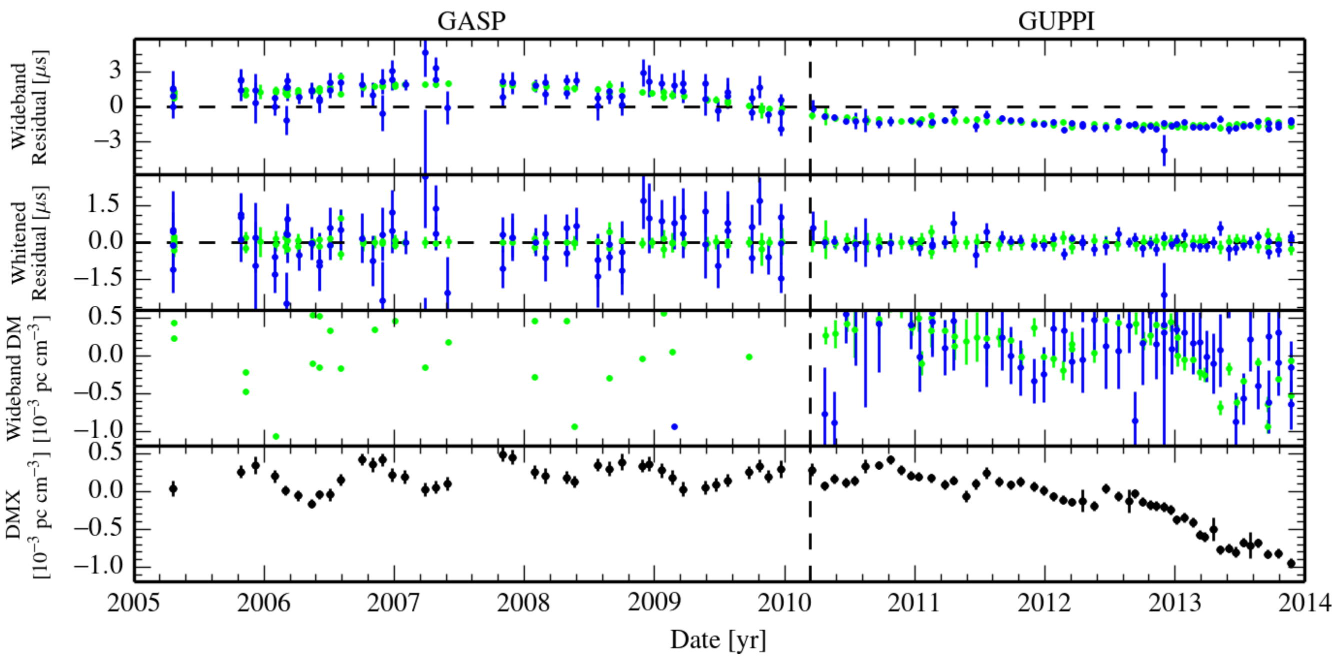

Fig. 4.51 – Timing summary panels for J0613−0200.

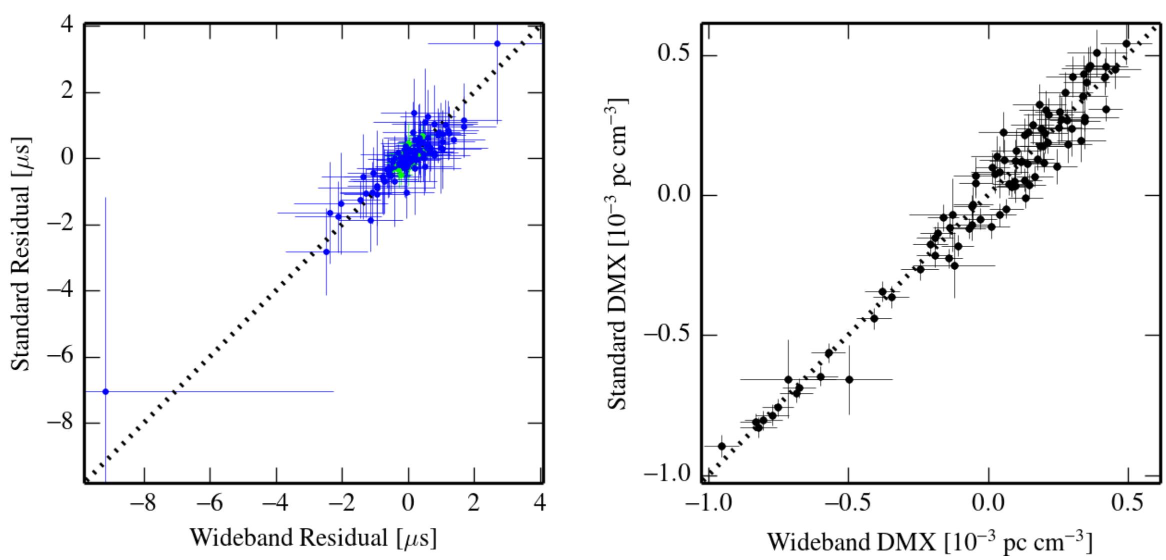

Fig. 4.52 – Comparison subplots for J0613−0200.

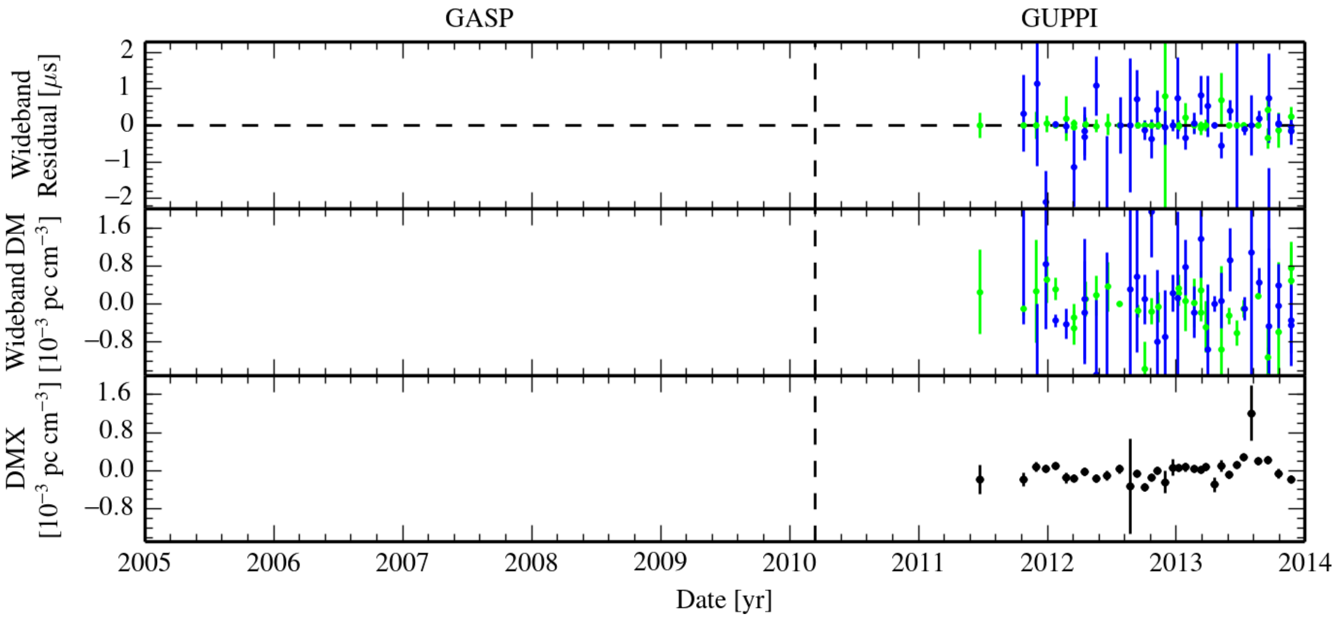

Fig. 4.53 – Timing summary panels for J0645+5158.

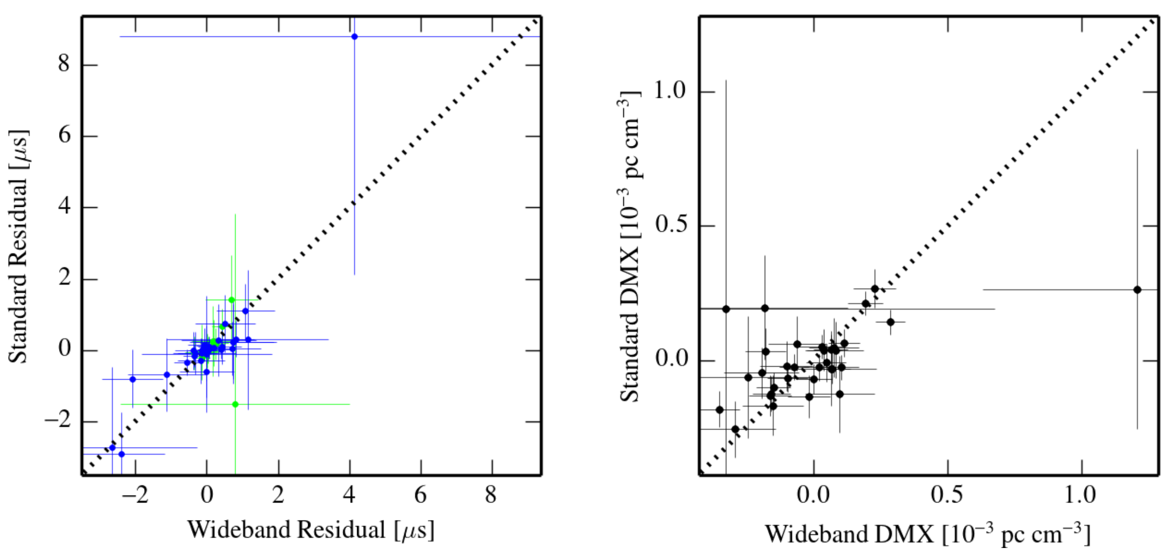

Fig. 4.54 – Comparison subplots for J0645+5158.

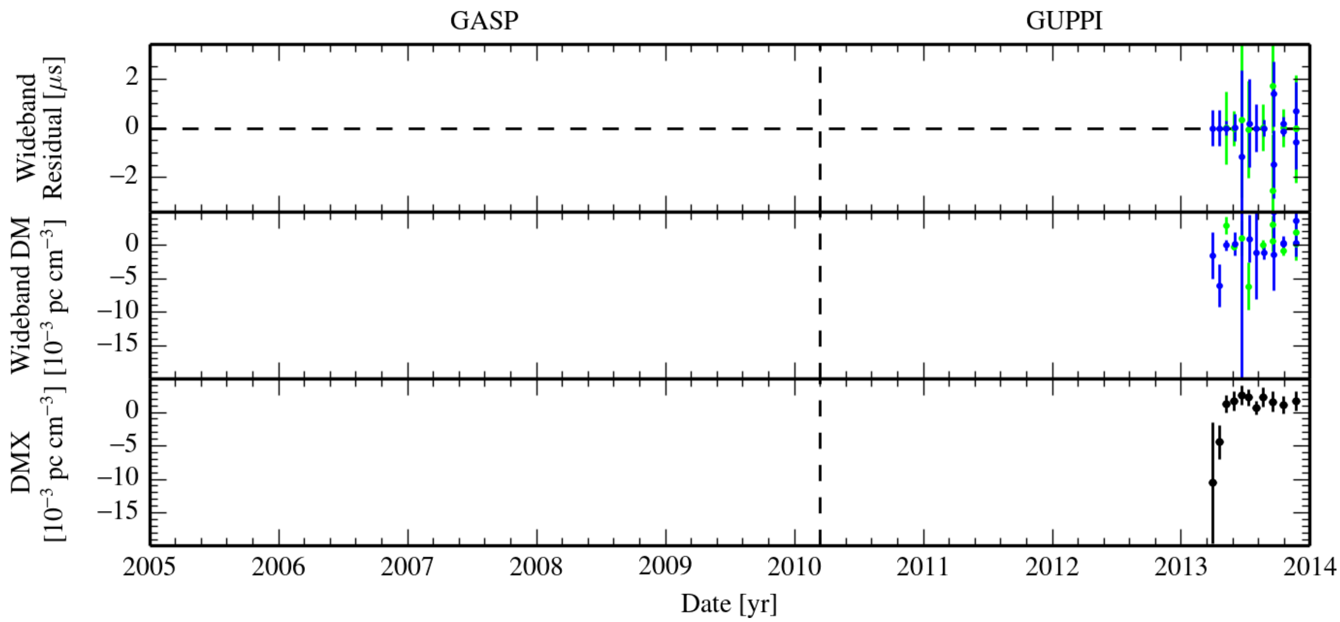

Fig. 4.55 – Timing summary panels for J0931−1902.

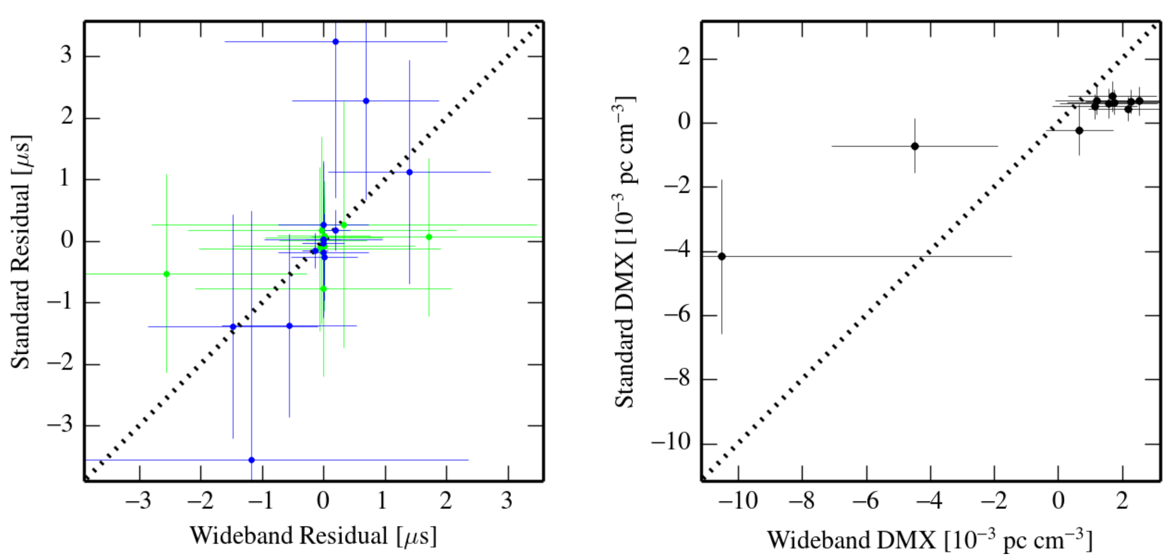

Fig. 4.56 – Comparison subplots for J0931−1902.

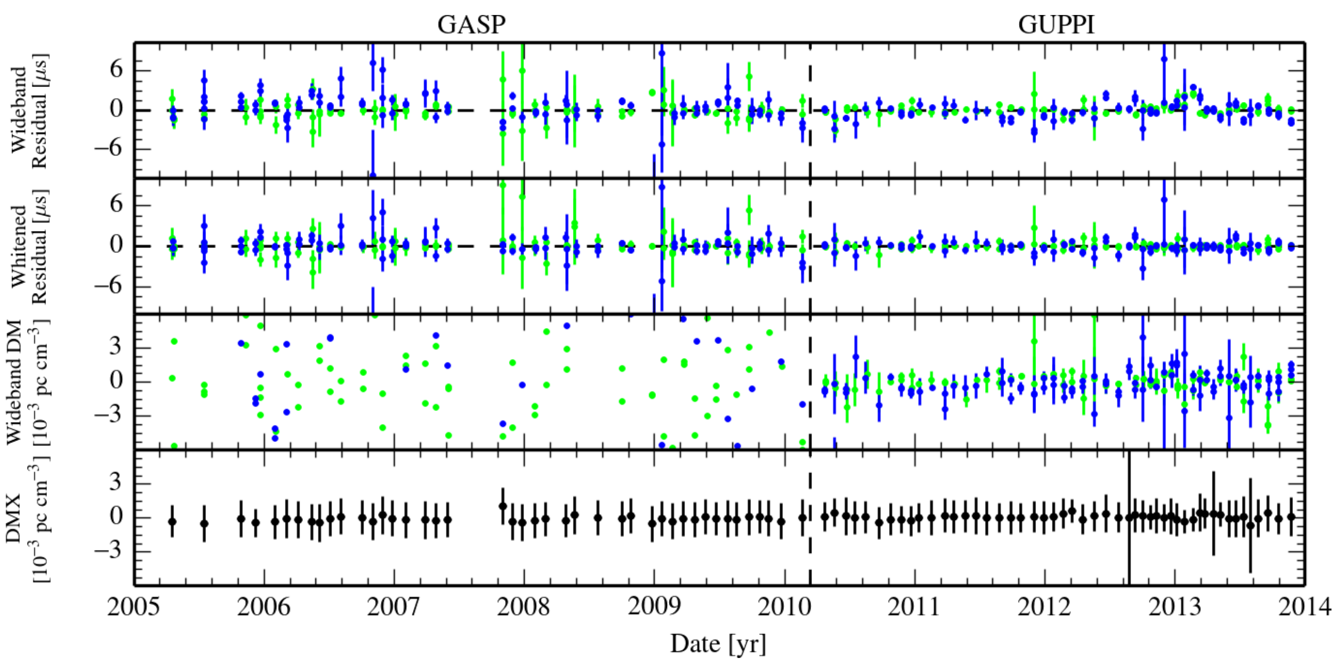

Fig. 4.57 – Timing summary panels for J1012+5307.

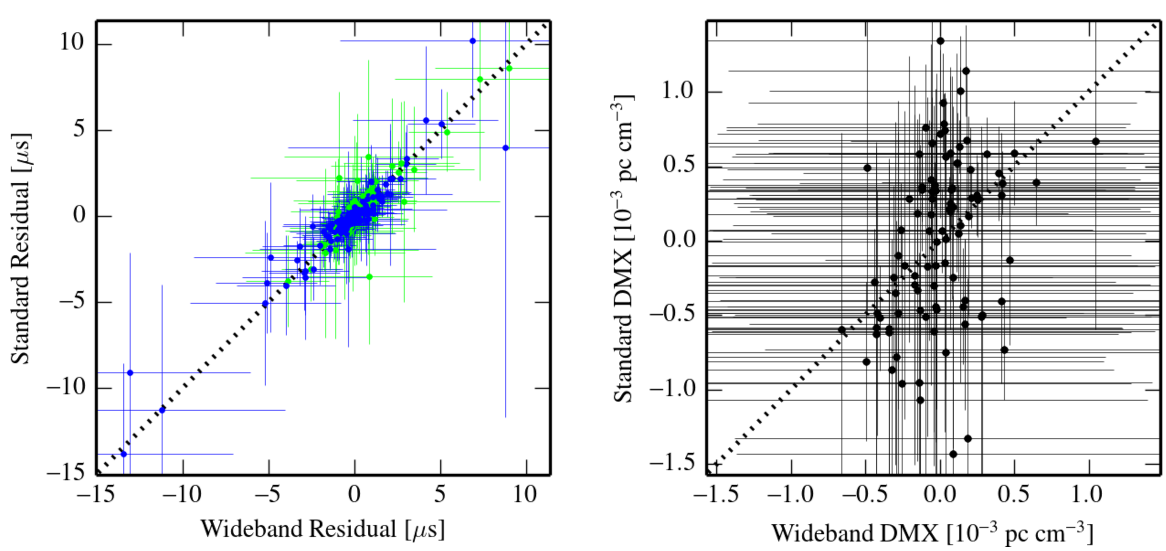

Fig. 4.58 – Comparison subplots for J1012+5307.

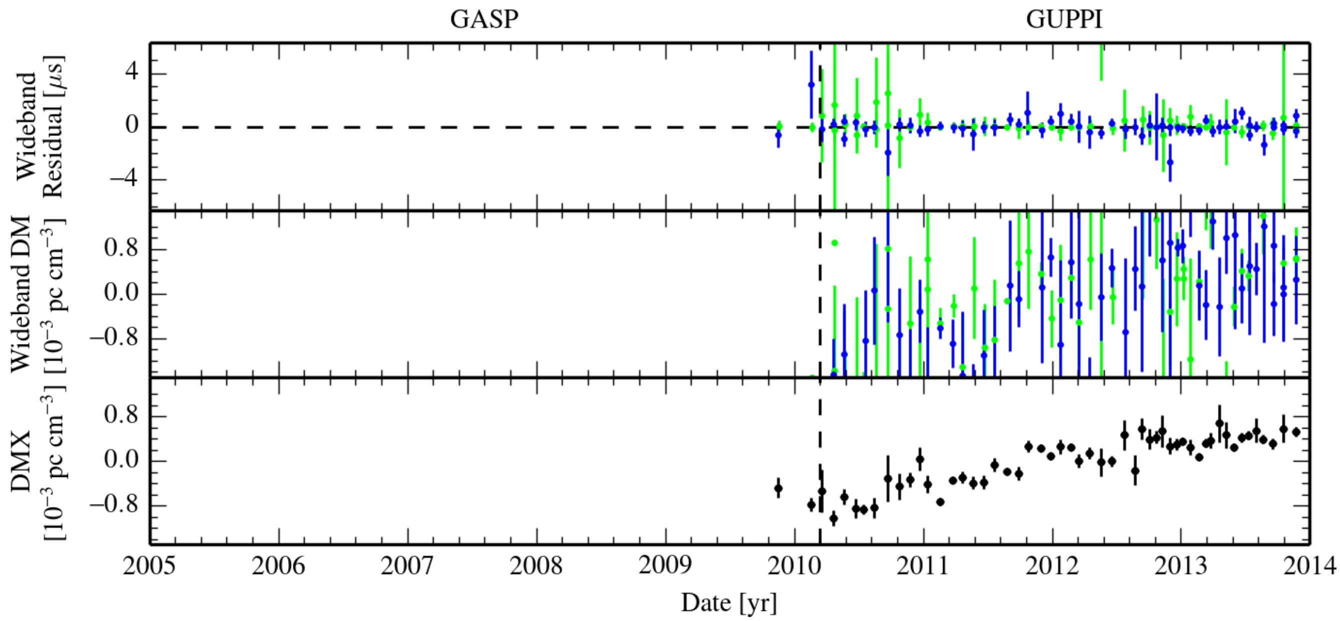

Fig. 4.59 – Timing summary panels for J1024−0719.

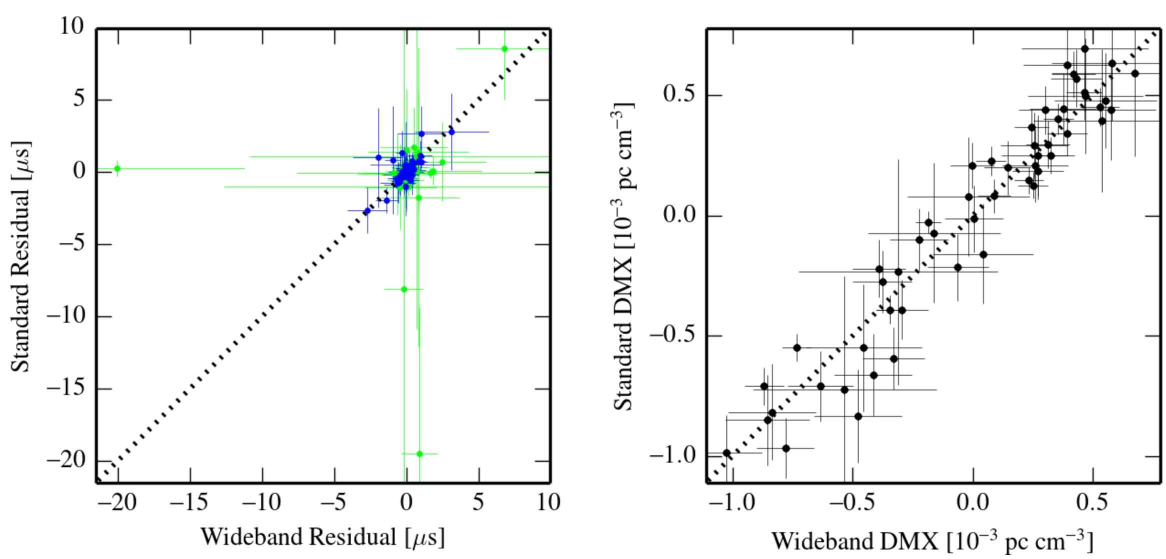

Fig. 4.60 – Comparison subplots for J1024−0719.

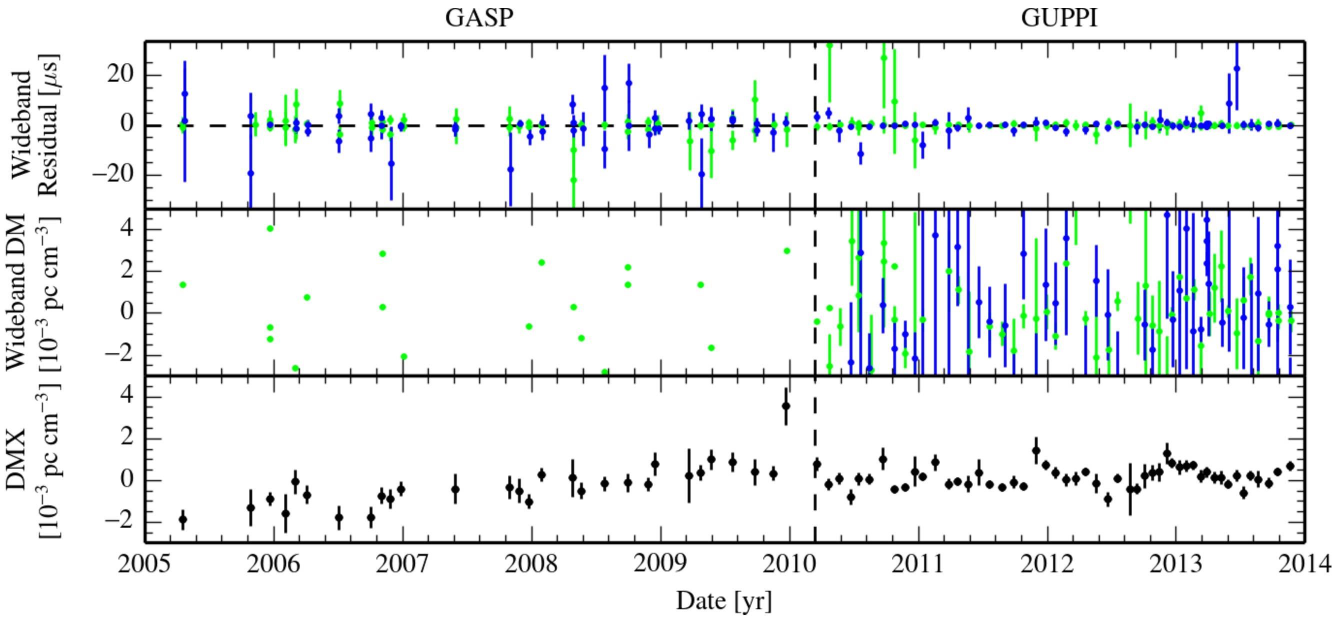

Fig. 4.61 – Timing summary panels for J1455−3330.

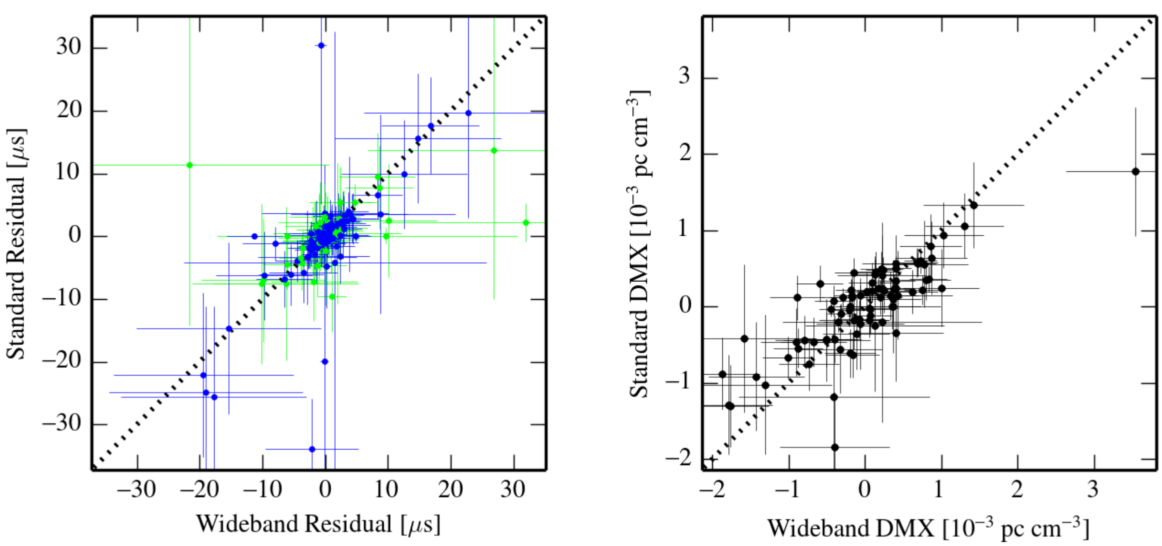

Fig. 4.62 – Comparison subplots for J1455−3330.

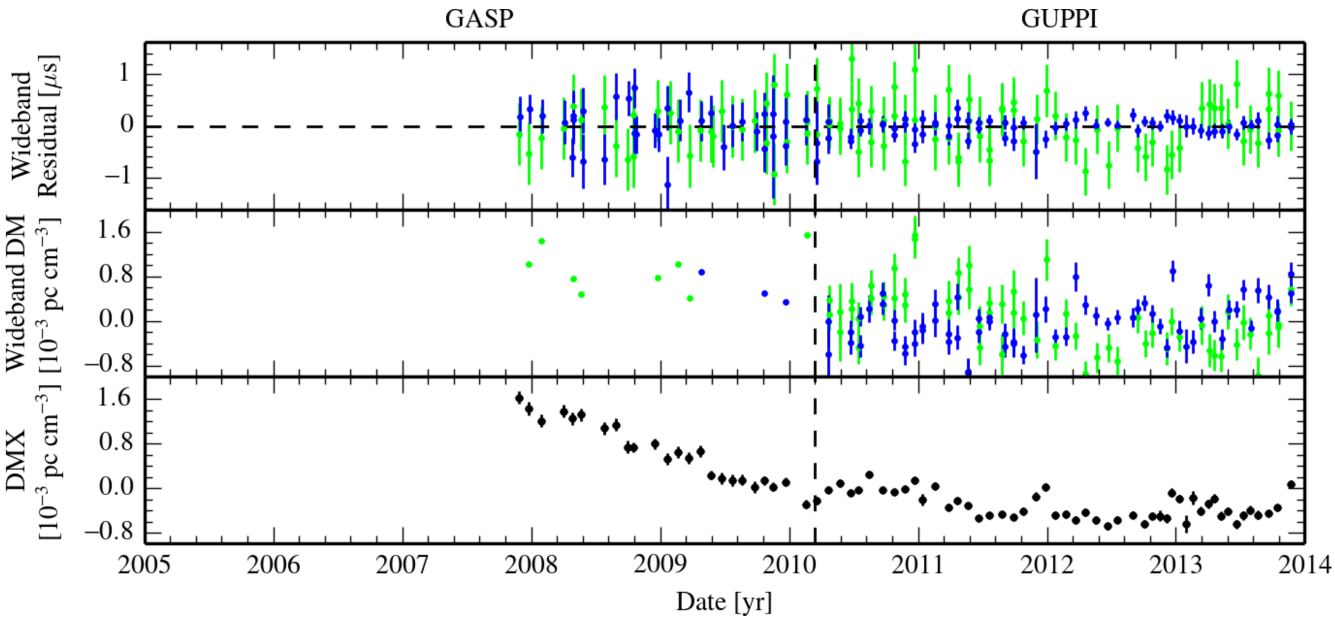

Fig. 4.63 – Timing summary panels for J1600−3053.

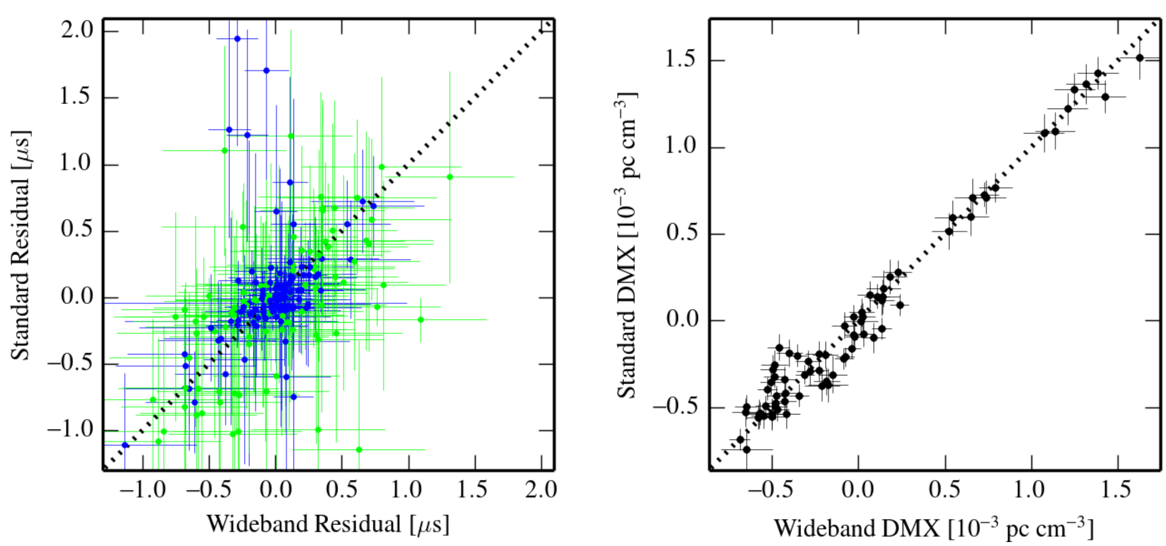

Fig. 4.64 – Comparison subplots for J1600−3053.

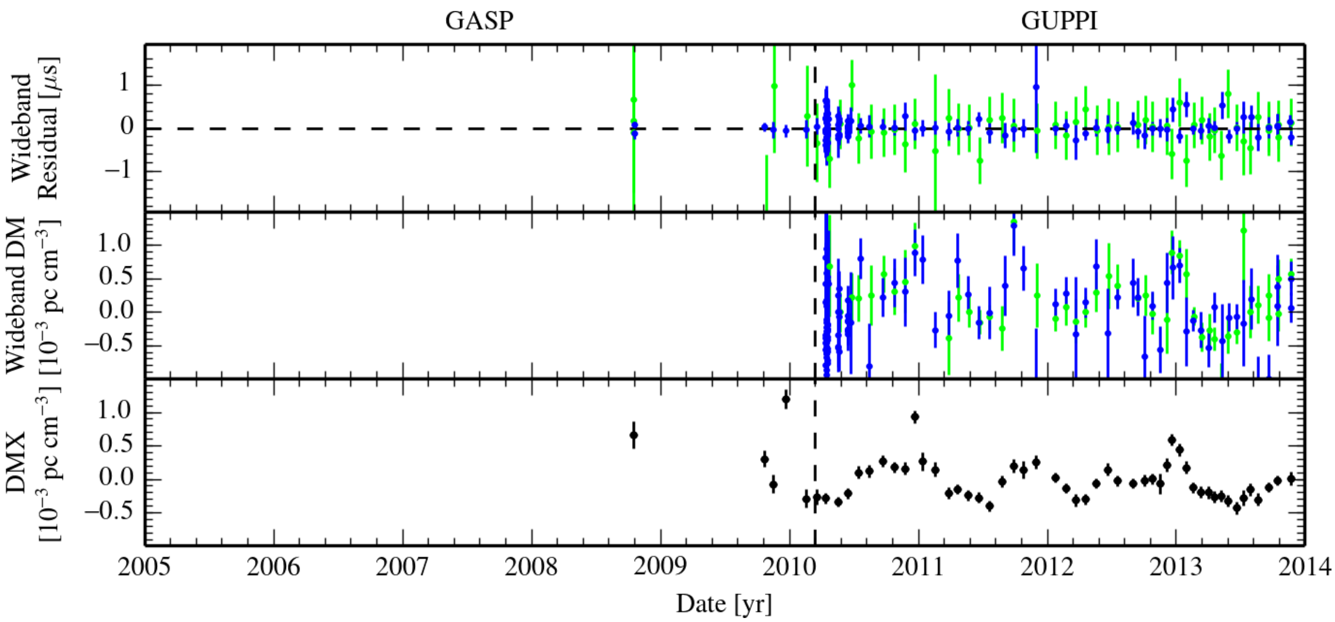

Fig. 4.65 – Timing summary panels for J1614−2230.

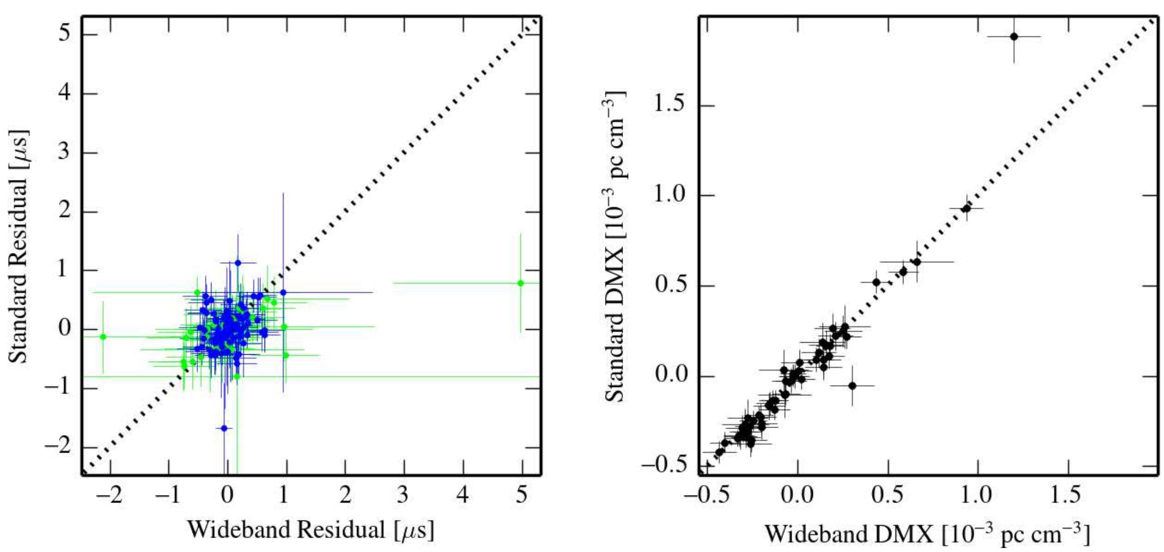

Fig. 4.66 – Comparison subplots for J1614−2230.

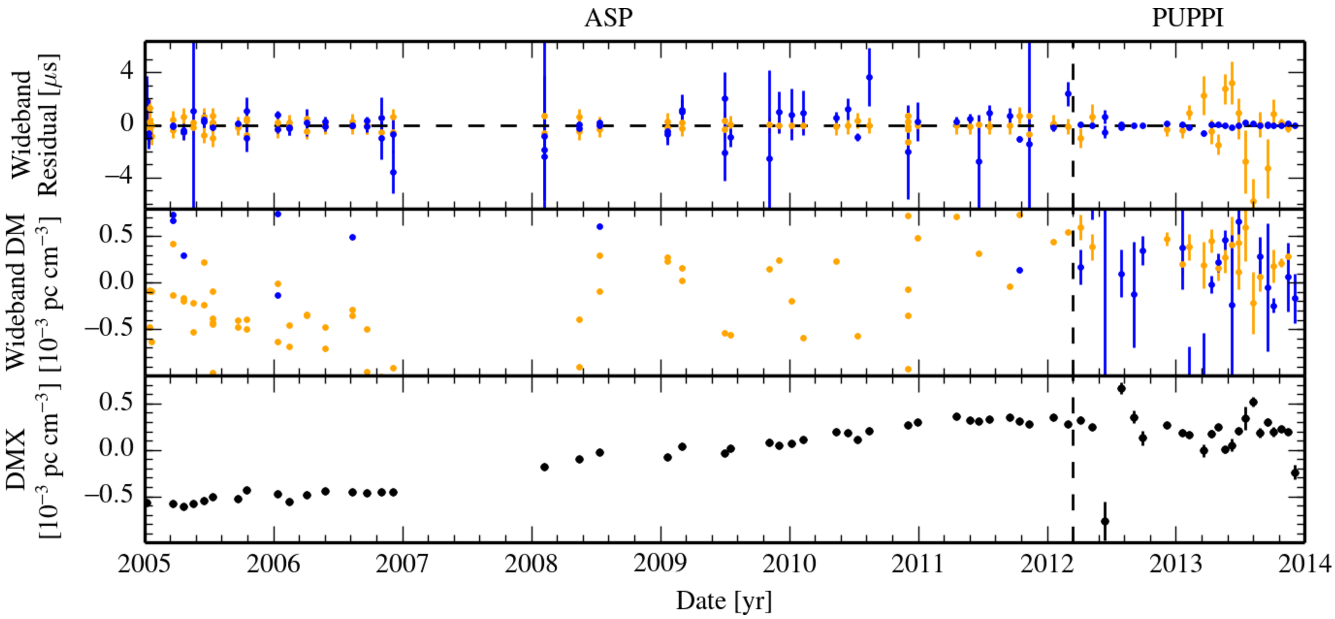

Fig. 4.67 – Timing summary panels for J1640+2224.

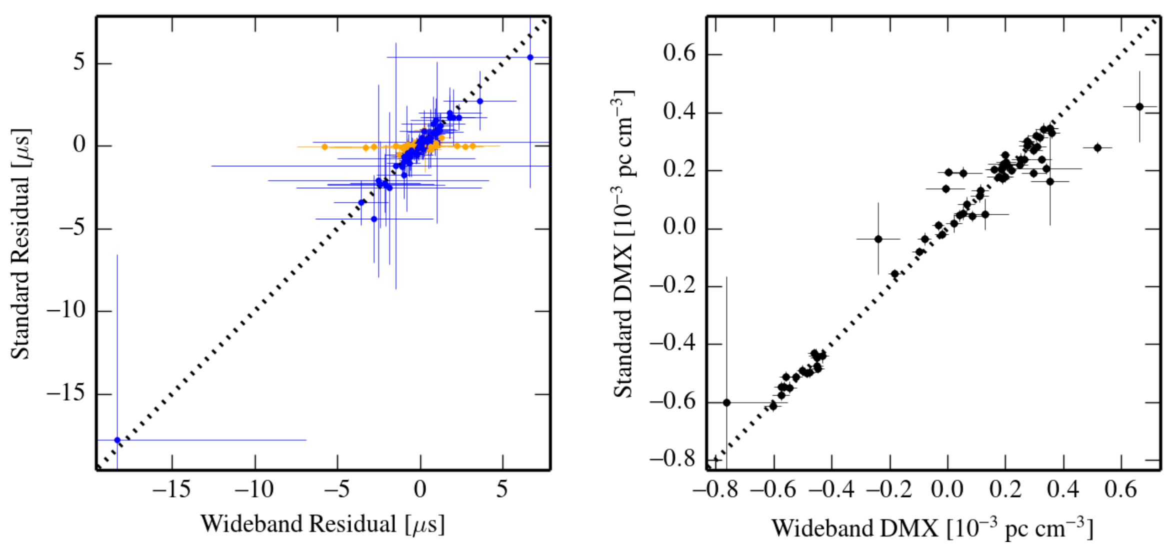

Fig. 4.68 – Comparison subplots for J1640+2224.

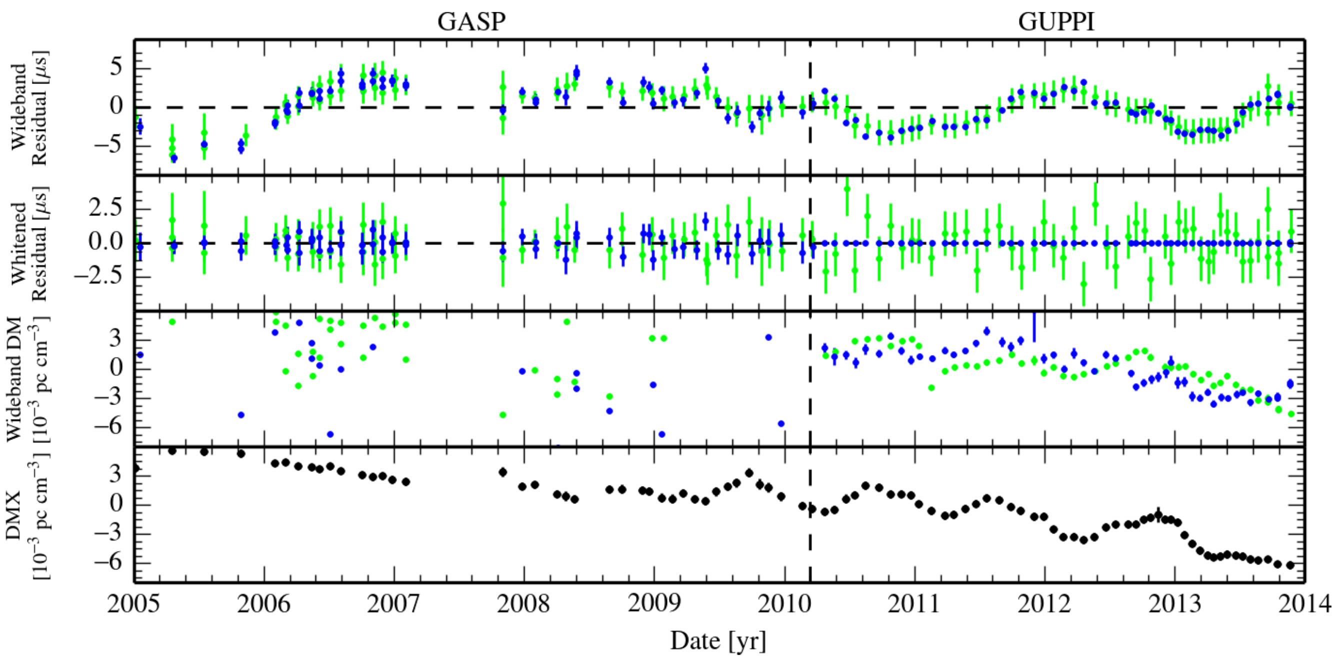

Fig. 4.69 – Timing summary panels for J1643−1224.

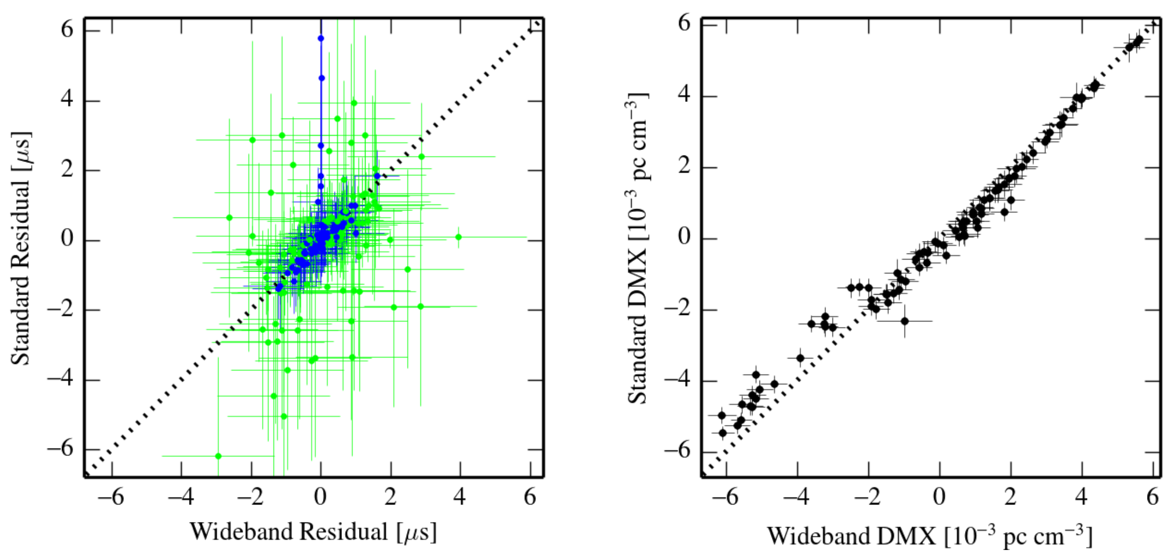

Fig. 4.70 – Comparison subplots for J1643−1224.

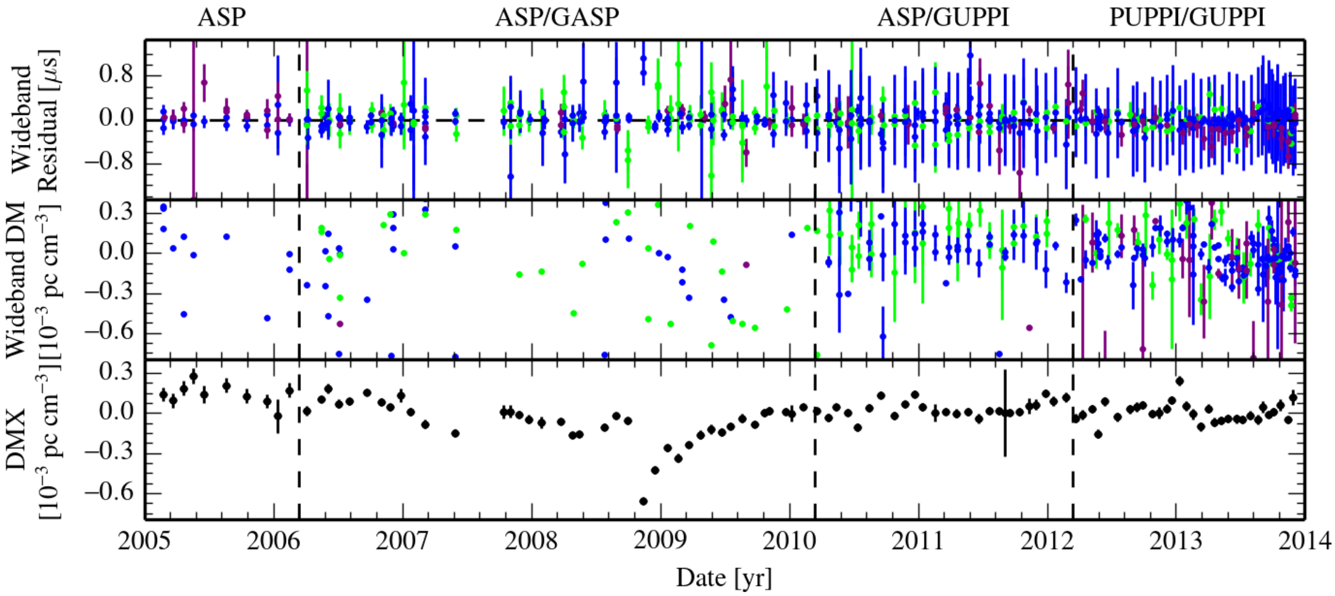

Fig. 4.71 – Timing summary panels for J1713+0747.

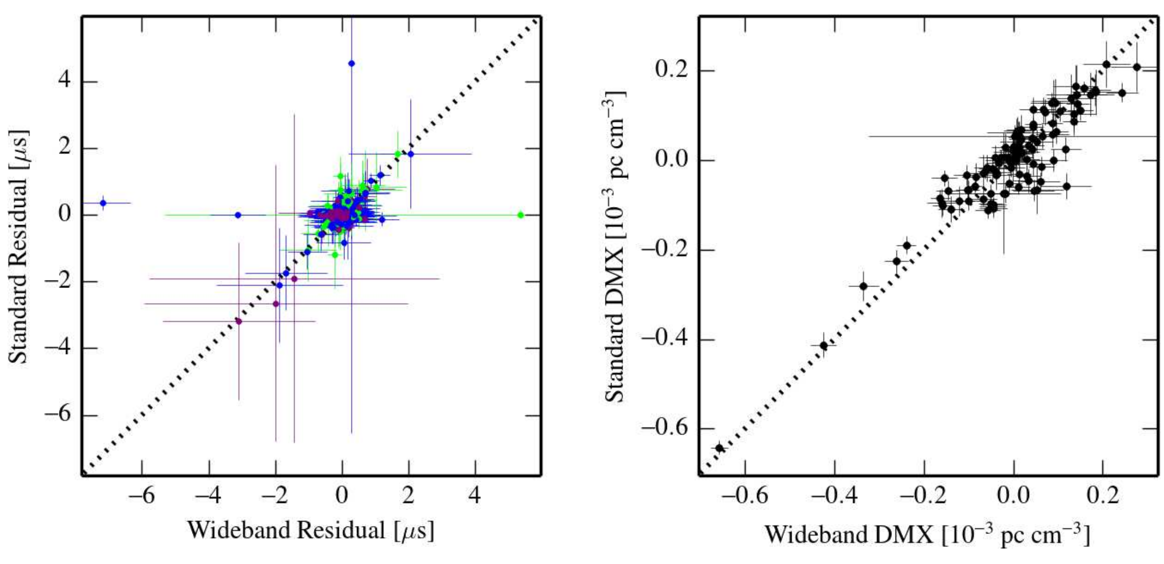

Fig. 4.72 – Comparison subplots for J1713+0747.

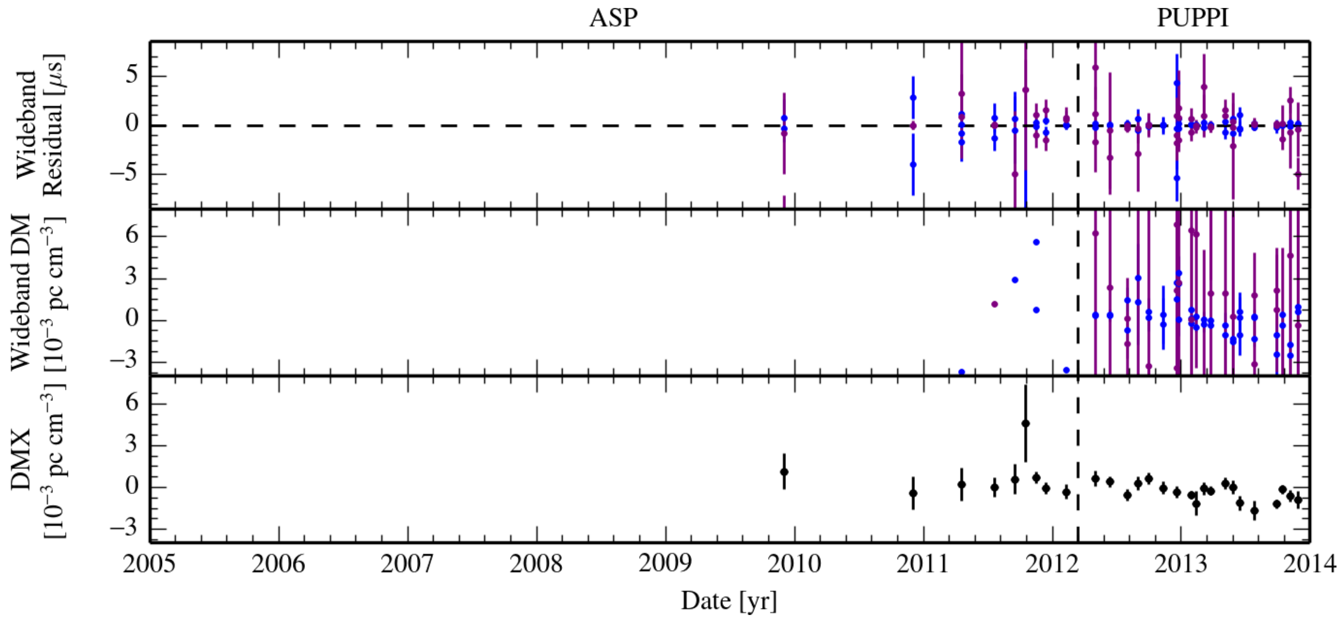

Fig. 4.73 – Timing summary panels for J1738+0333.

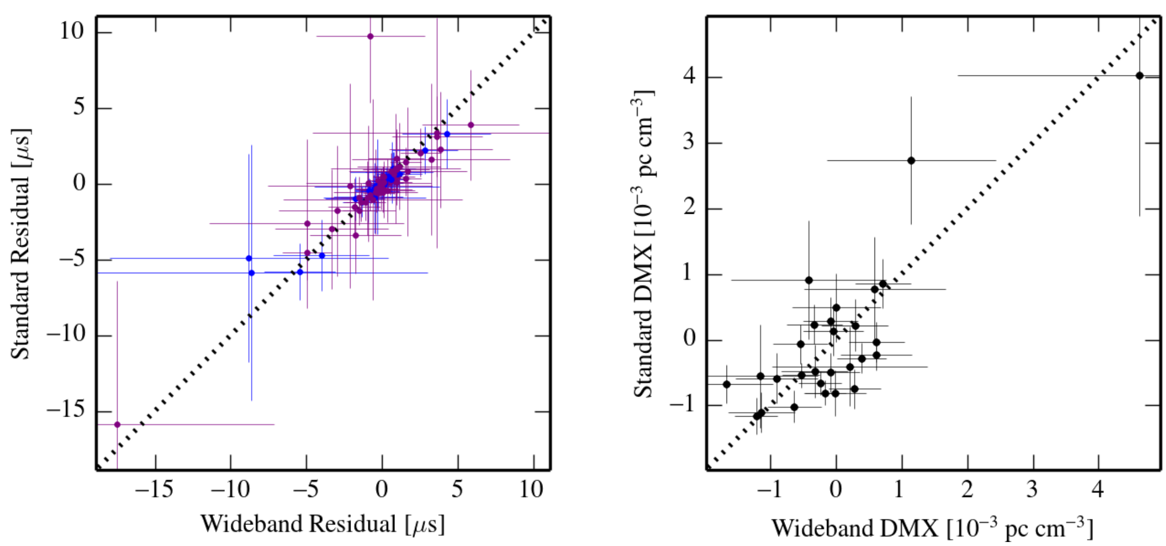

Fig. 4.74 – Comparison subplots for J1738+0333.

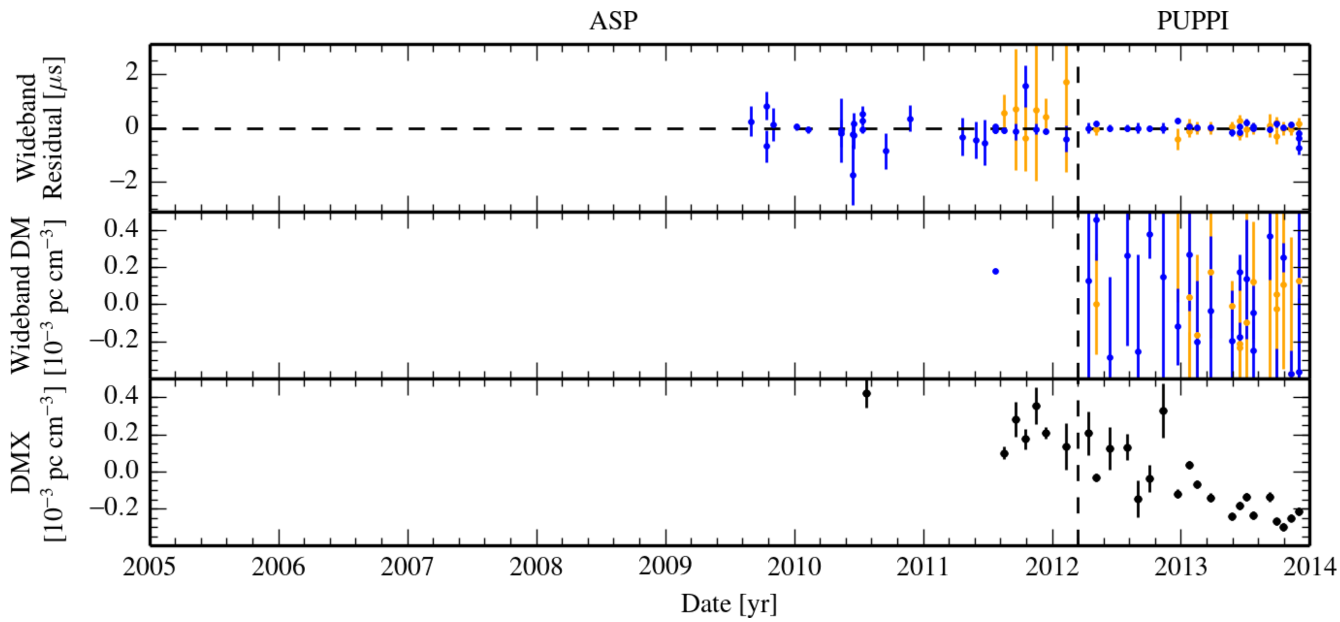

Fig. 4.75 – Timing summary panels for J1741+1351.

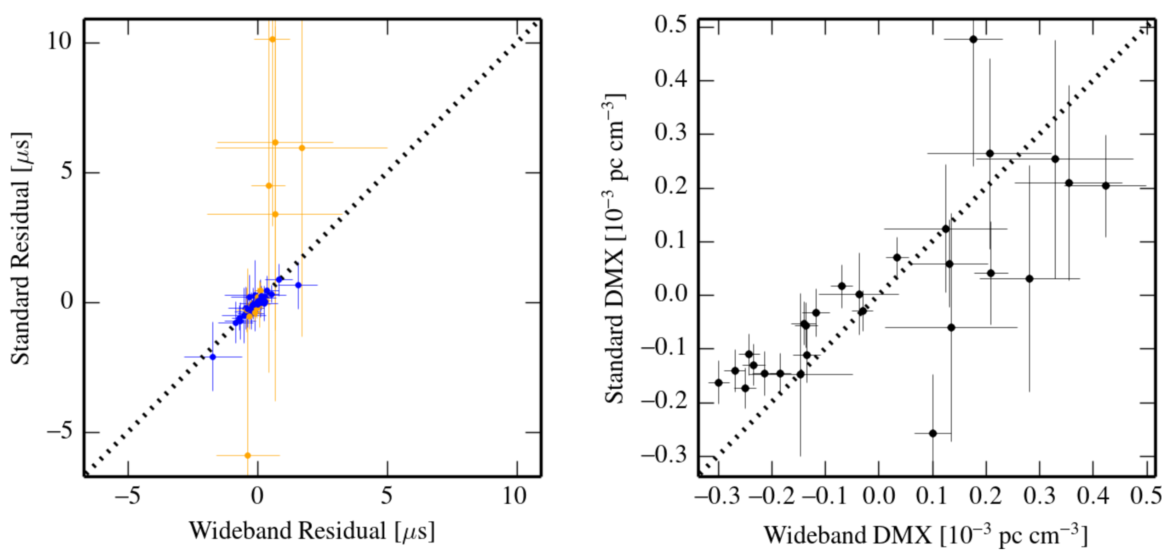

Fig. 4.76 – Comparison subplots for J1741+1351.

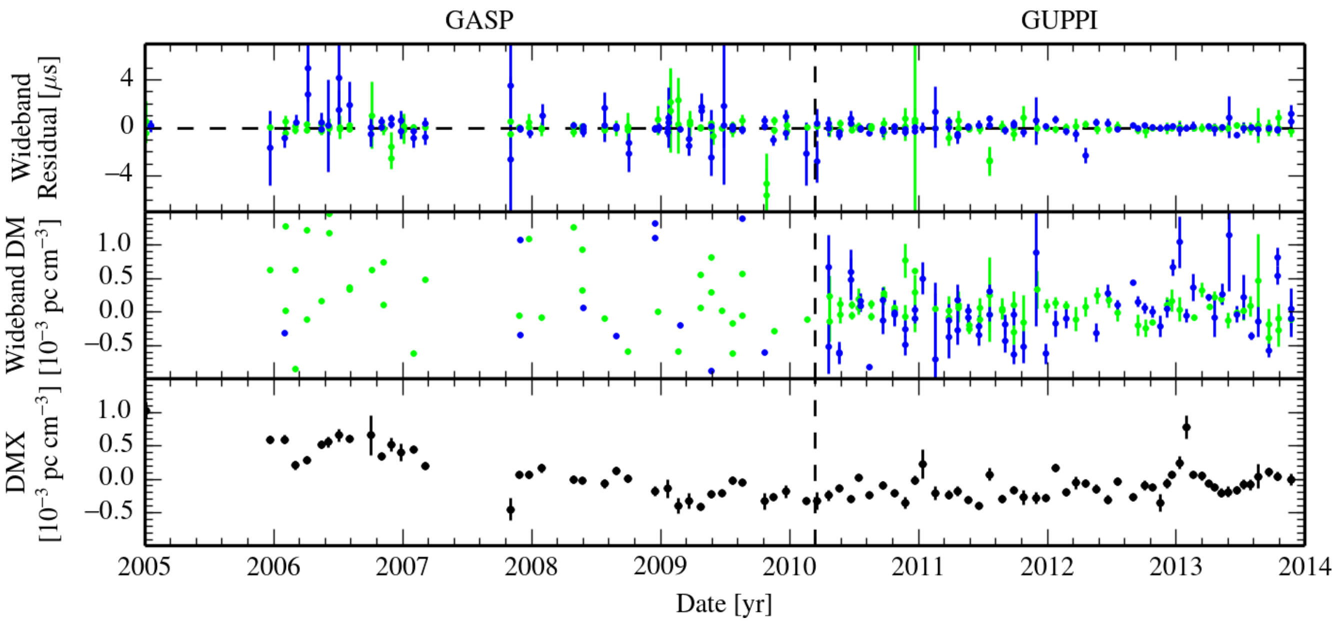

Fig. 4.77 – Timing summary panels for J1744−1134.

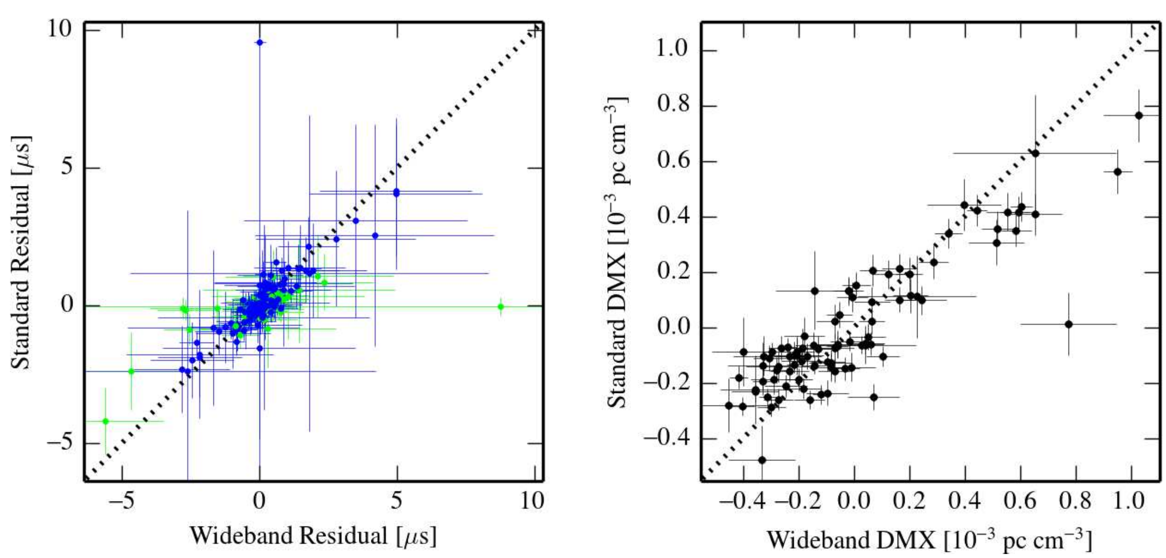

Fig. 4.78 – Comparison subplots for J1744−1134.

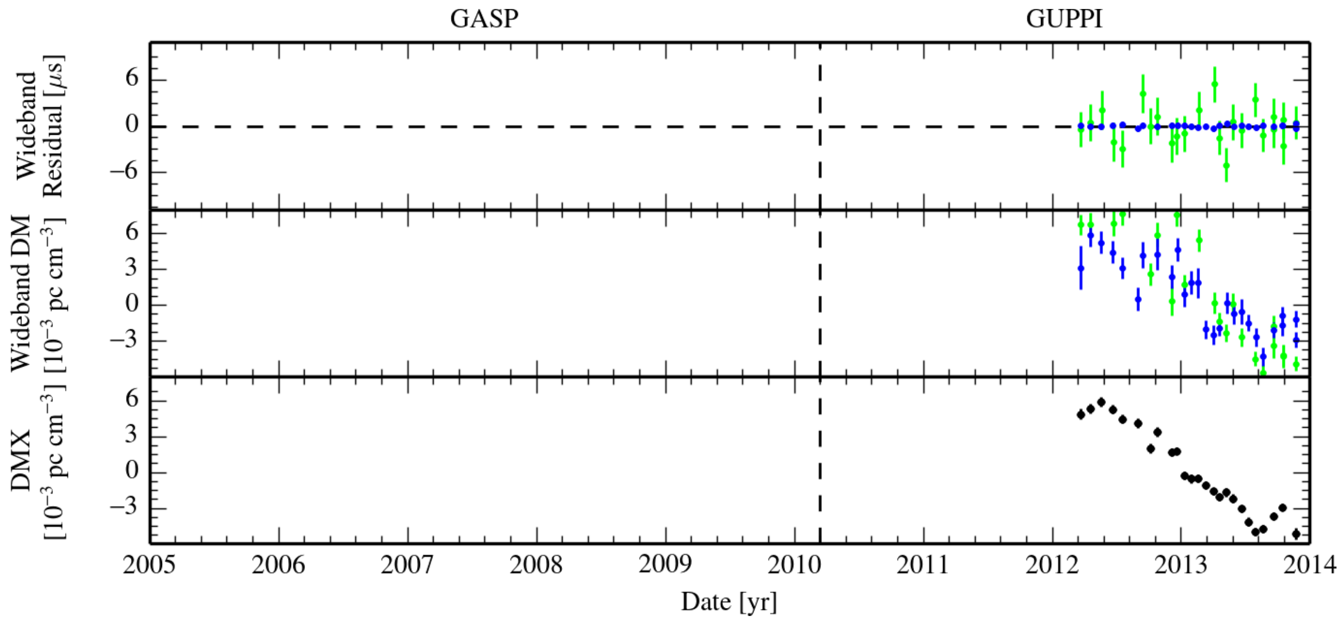

Fig. 4.79 – Timing summary panels for J1747−4036.

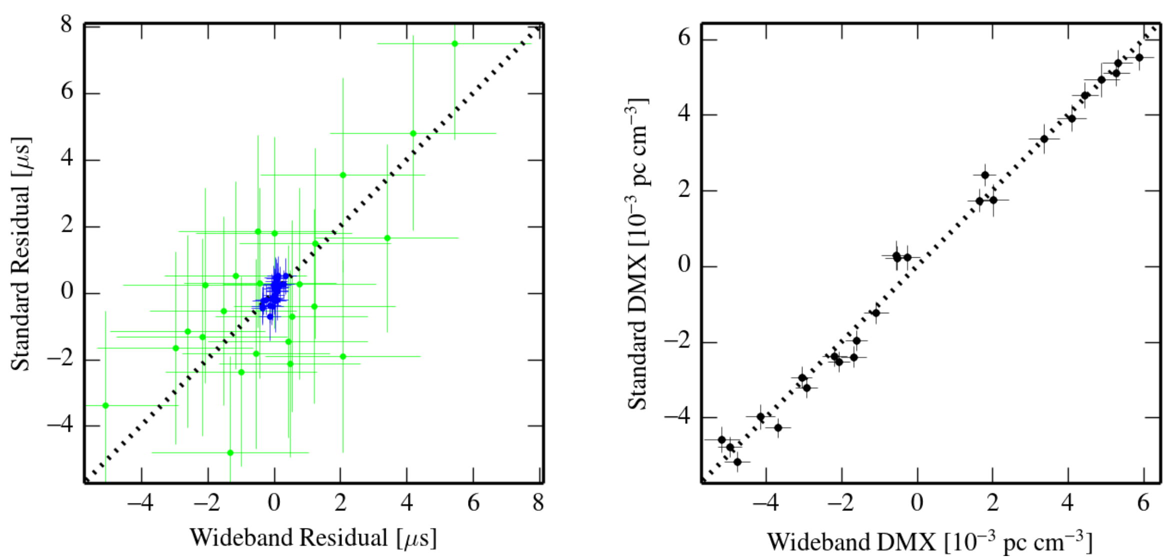

Fig. 4.80 – Comparison subplots for J1747−4036.

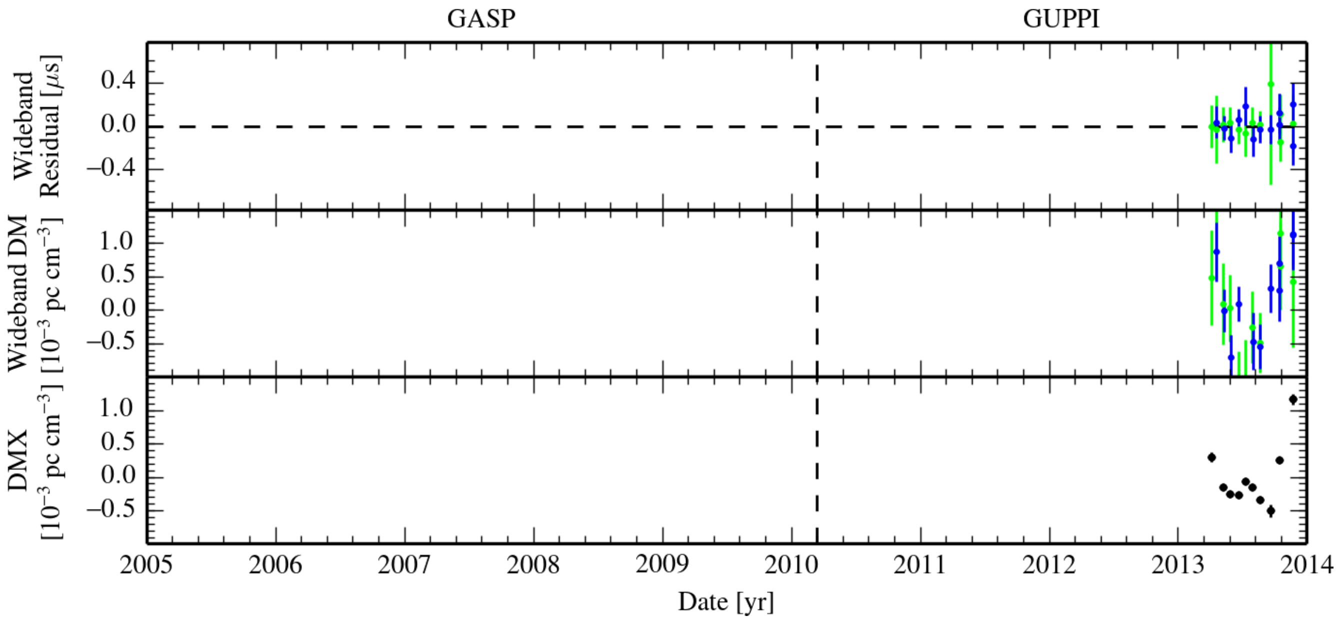

Fig. 4.81 – Timing summary panels for J1832−0836.

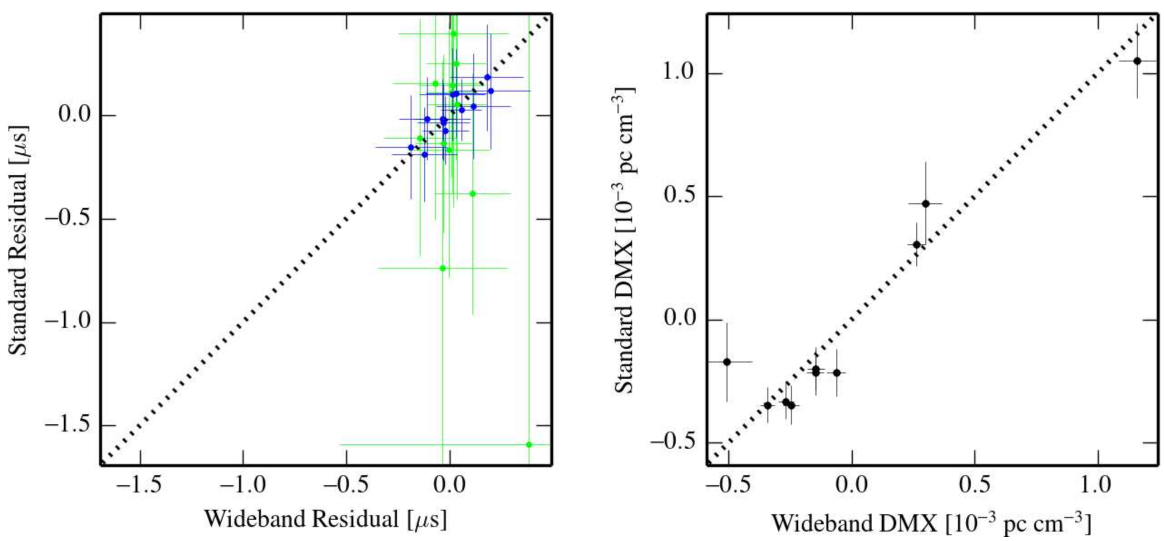

Fig. 4.82 – Comparison subplots for J1832−0836.

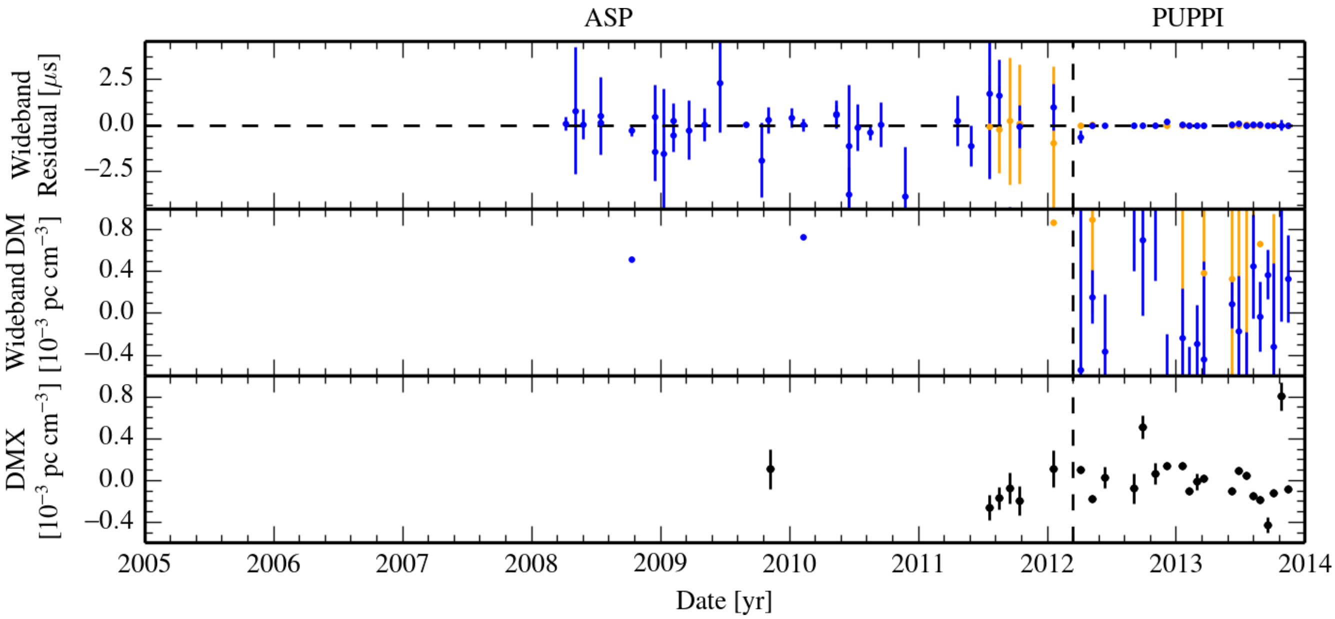

Fig. 4.83 – Timing summary panels for J1853+1303.

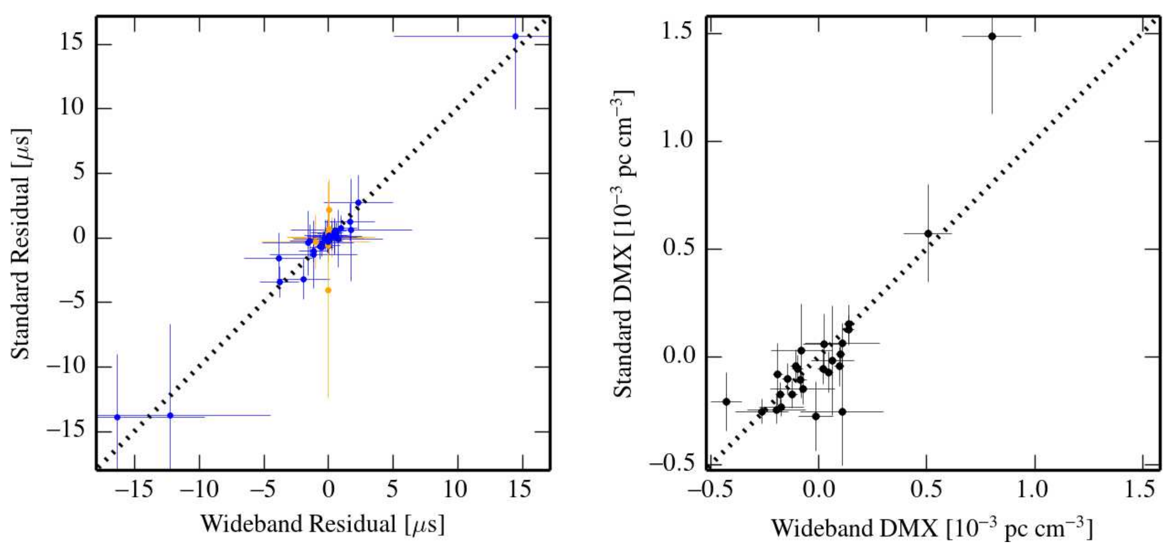

Fig. 4.84 – Comparison subplots for J1853+1303.

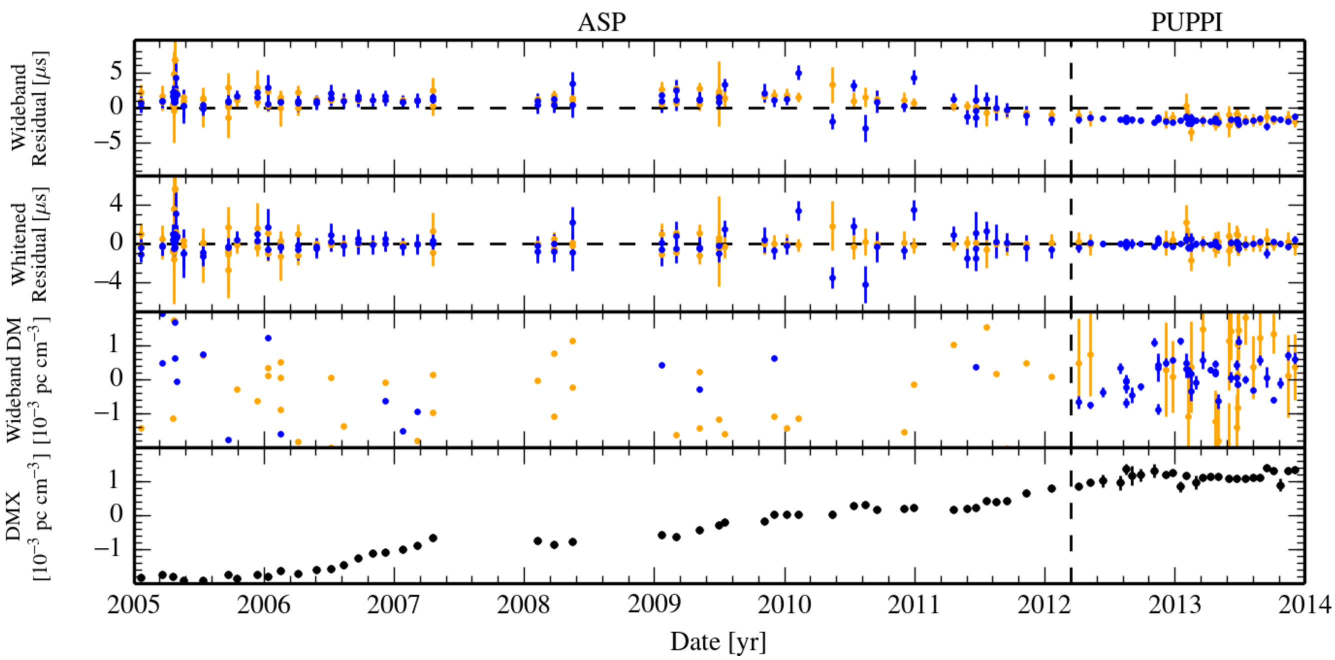

Fig. 4.85 – Timing summary panels for B1855+09.

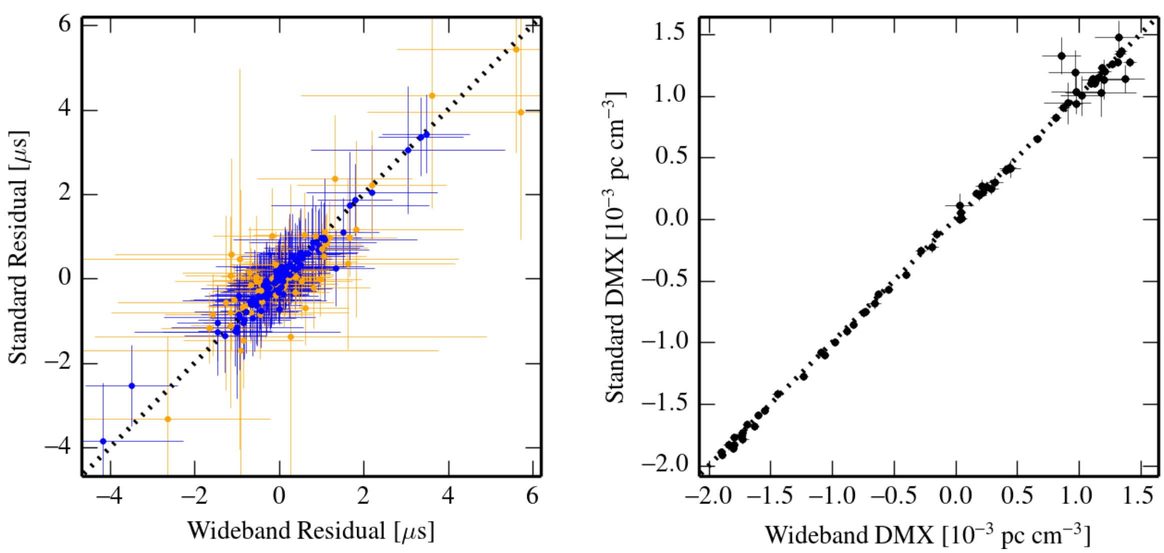

Fig. 4.86 – Comparison subplots for B1855+09.

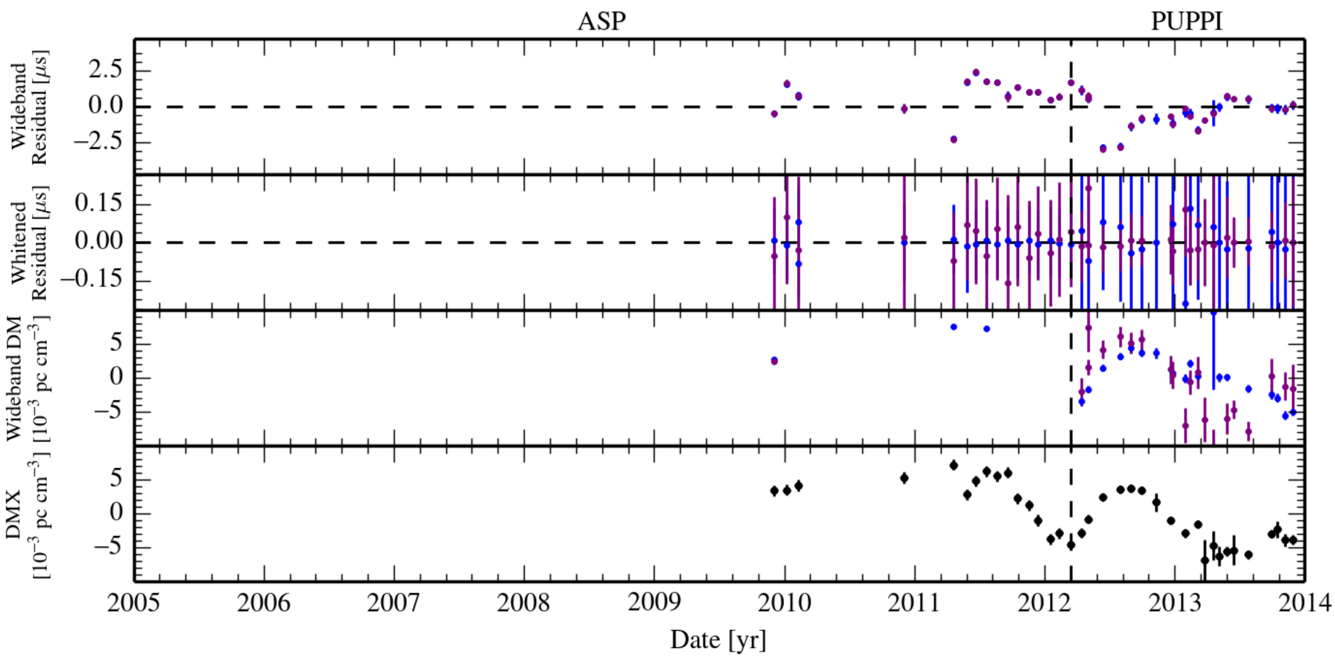

Fig. 4.87 – Timing summary panels for J1903+0327.

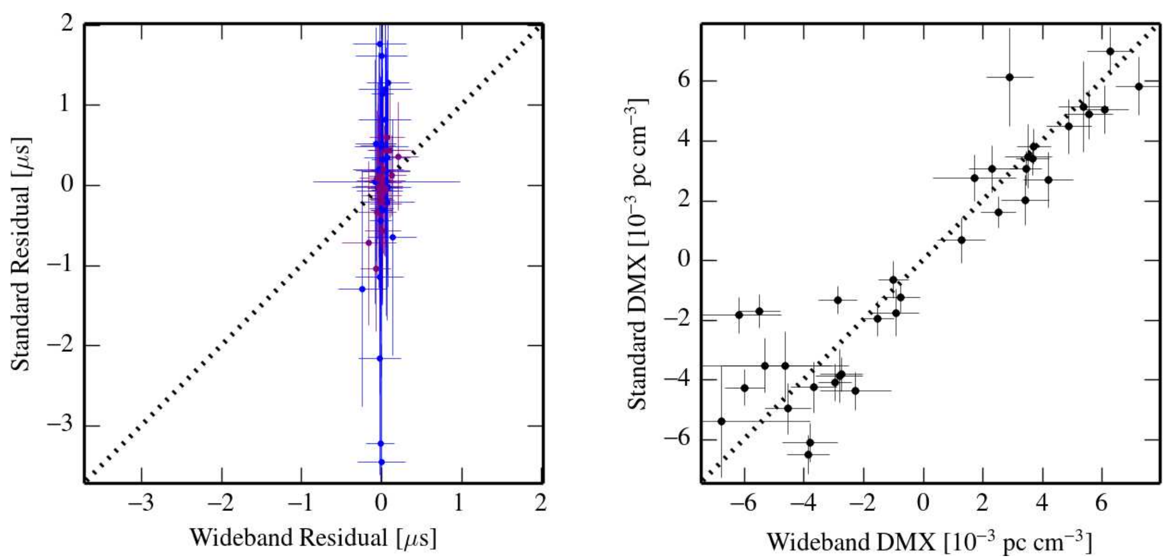

Fig. 4.88 – Comparison subplots for J1903+0327.

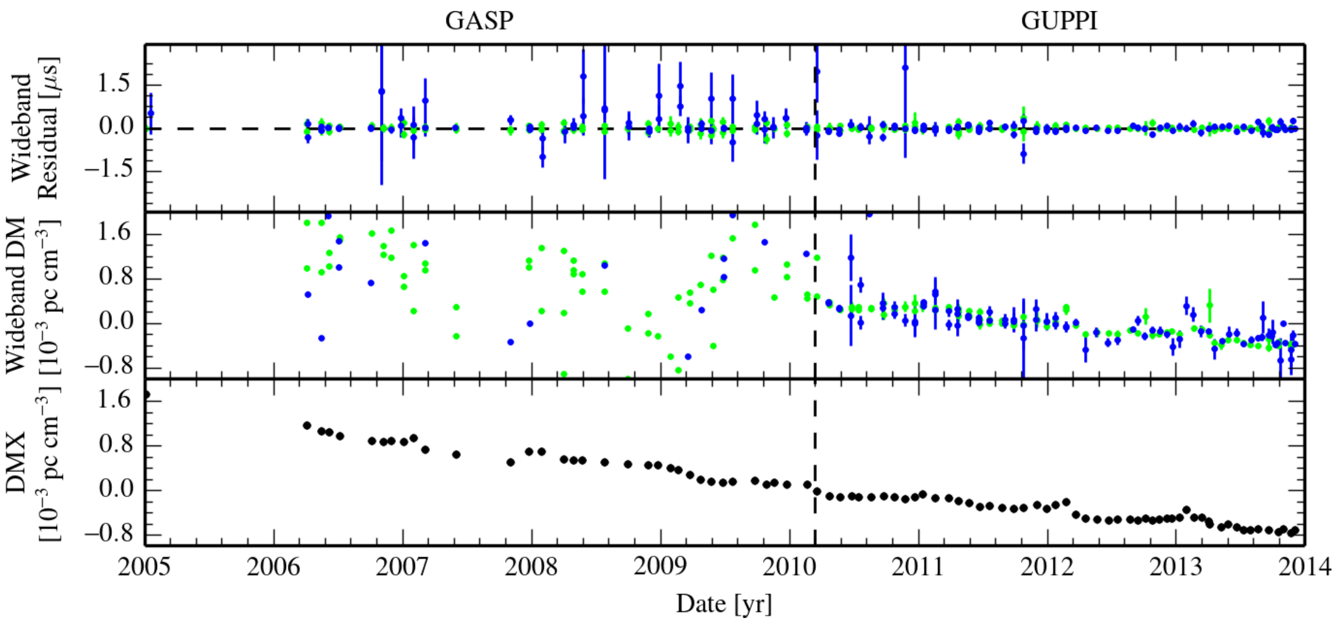

Fig. 4.89 – Timing summary panels for J1909−3744.

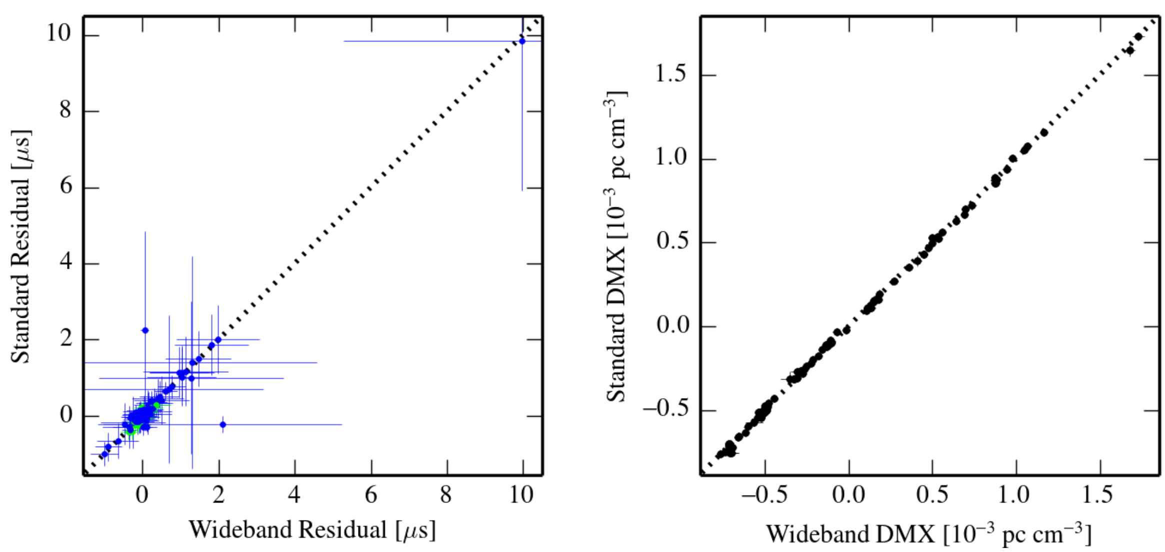

Fig. 4.90 – Comparison subplots for J1909−3744.

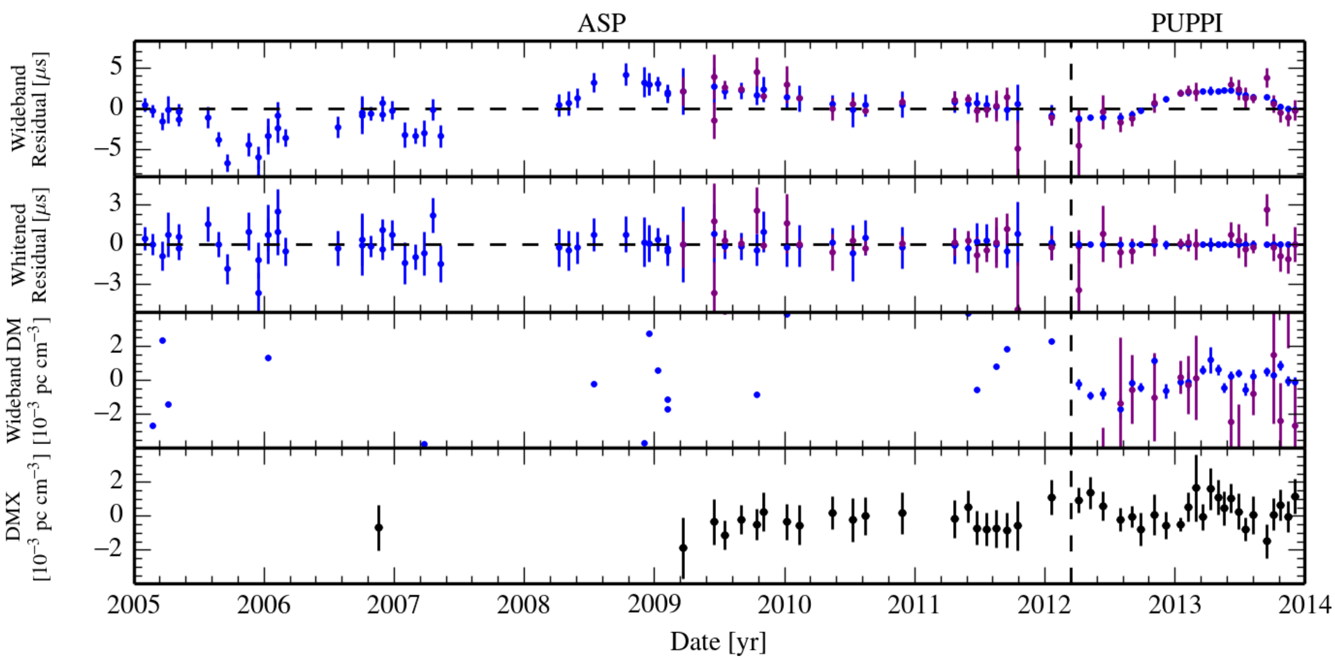

Fig. 4.91 – Timing summary panels for J1910+1256.

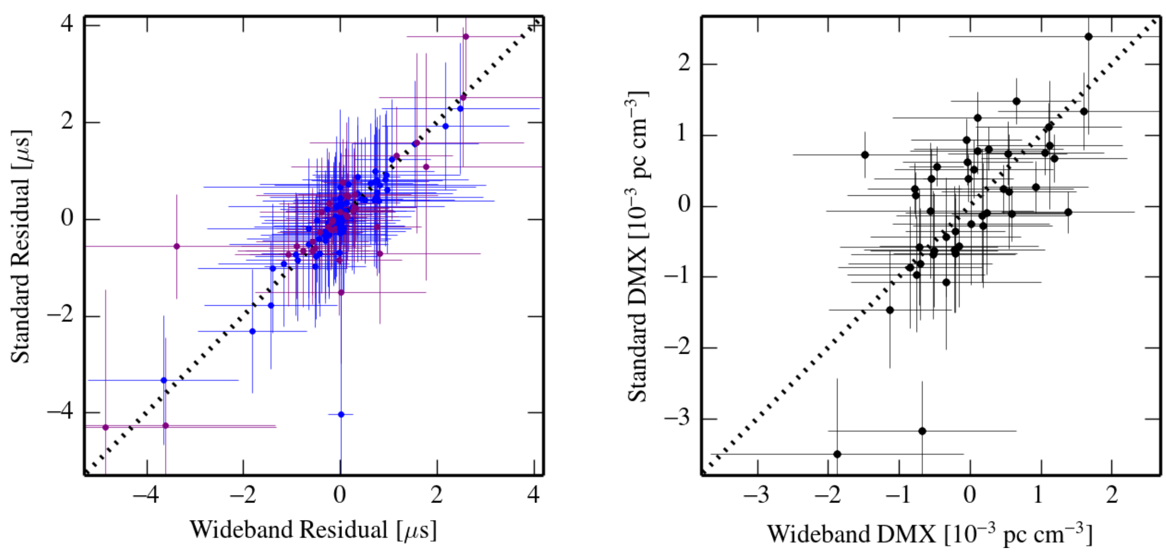

Fig. 4.92 – Comparison subplots for J1910+1256.

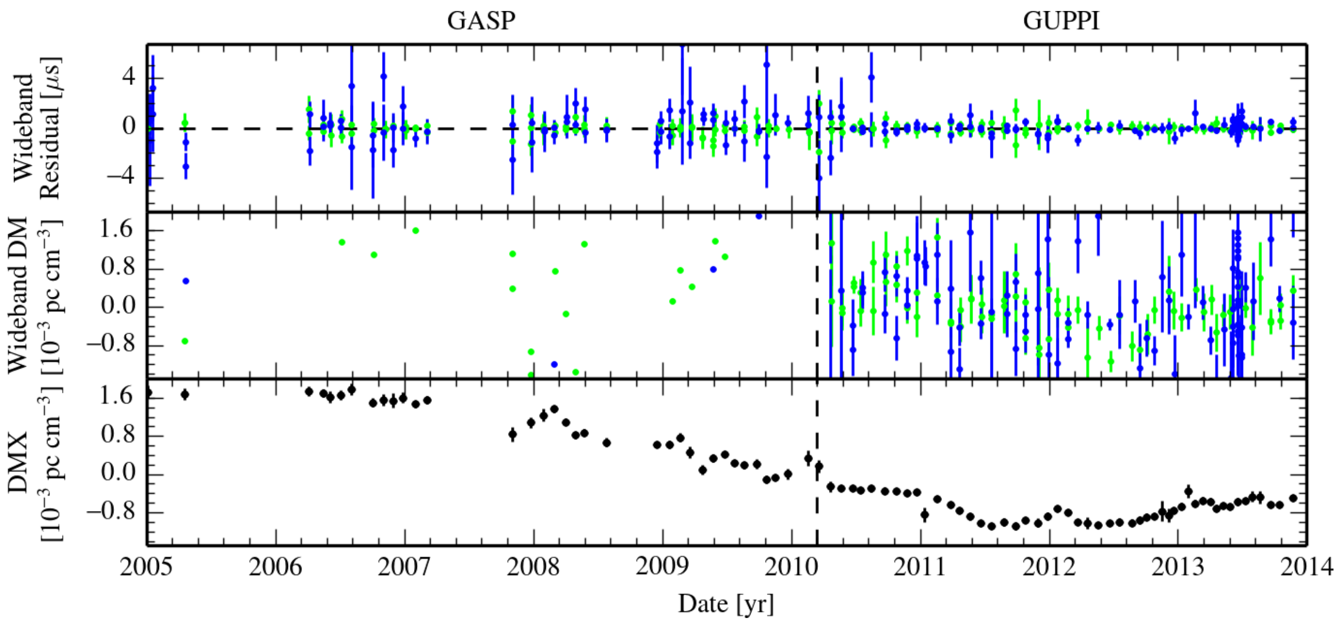

Fig. 4.93 – Timing summary panels for J1918−0642.

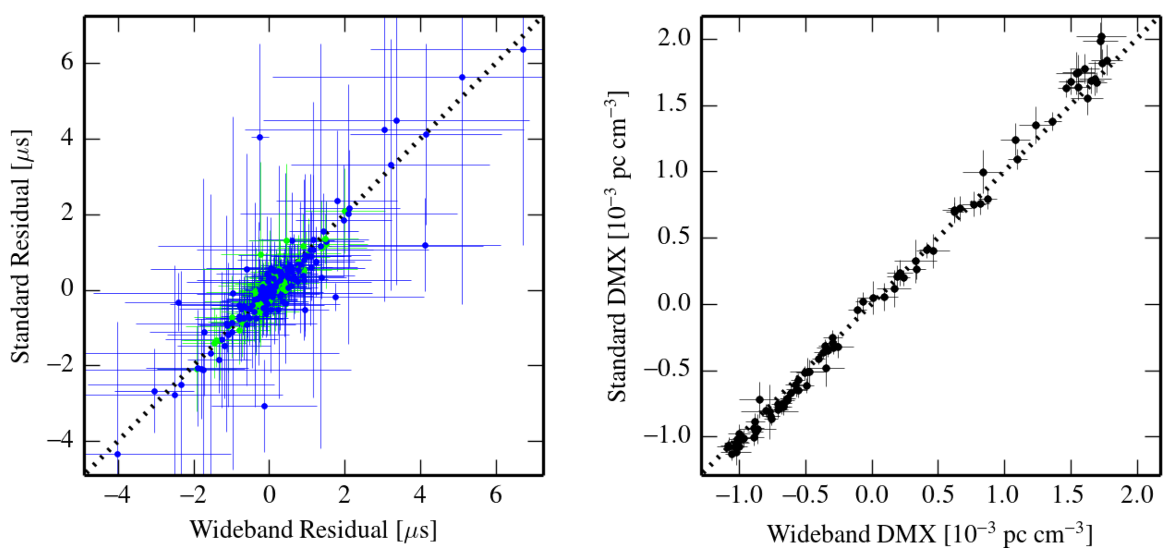

Fig. 4.94 – Comparison subplots for J1918−0642.

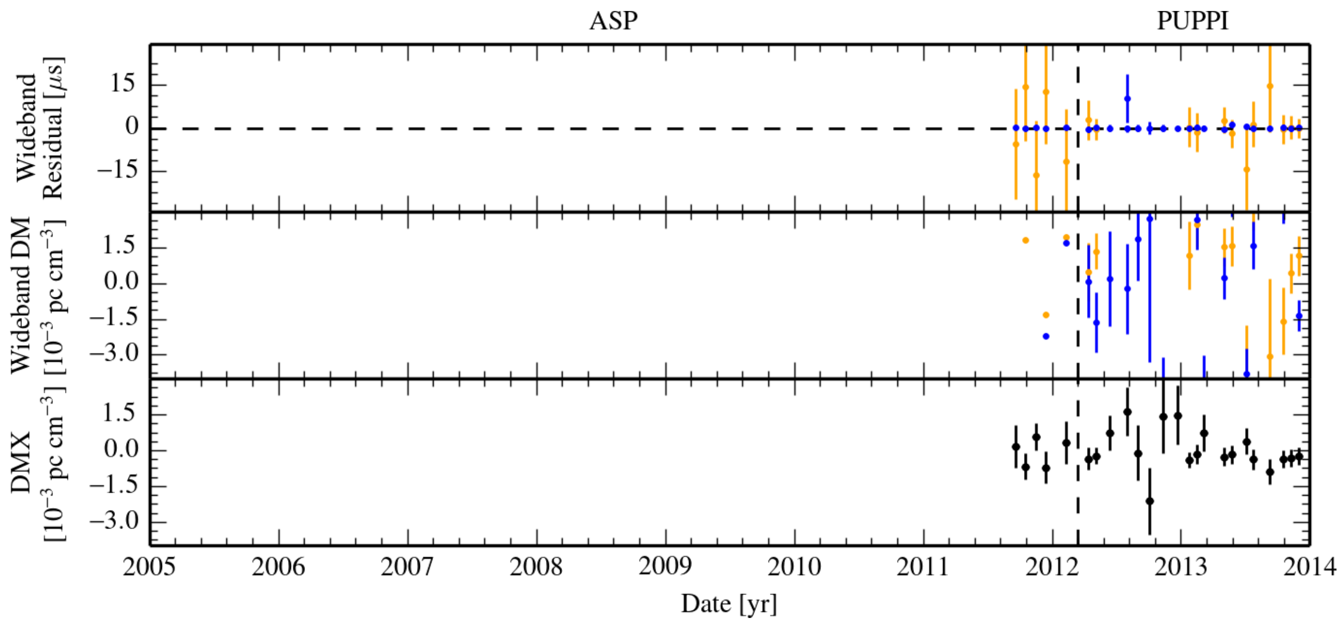

Fig. 4.95 – Timing summary panels for J1923+2515.

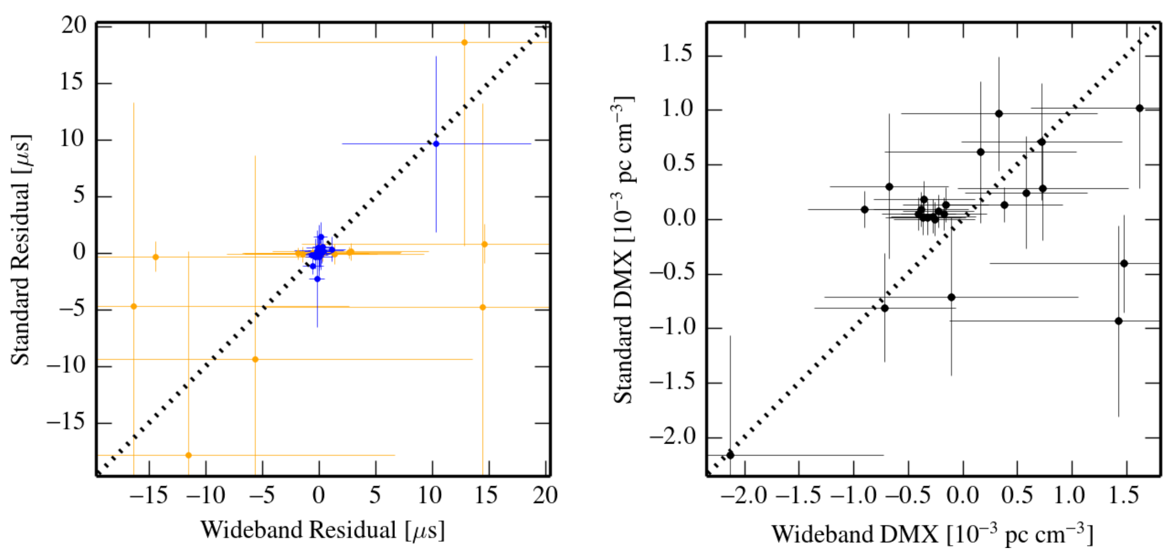

Fig. 4.96 – Comparison subplots for J1923+2515.

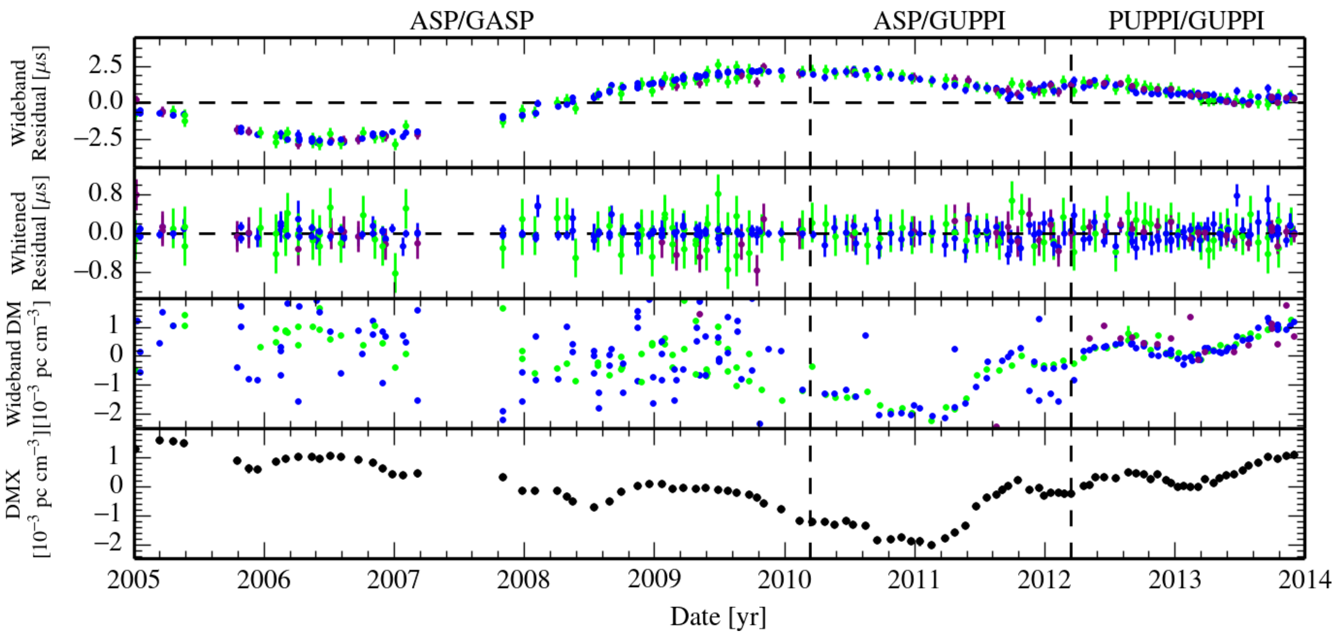

Fig. 4.97 – Timing summary panels for B1937+21.

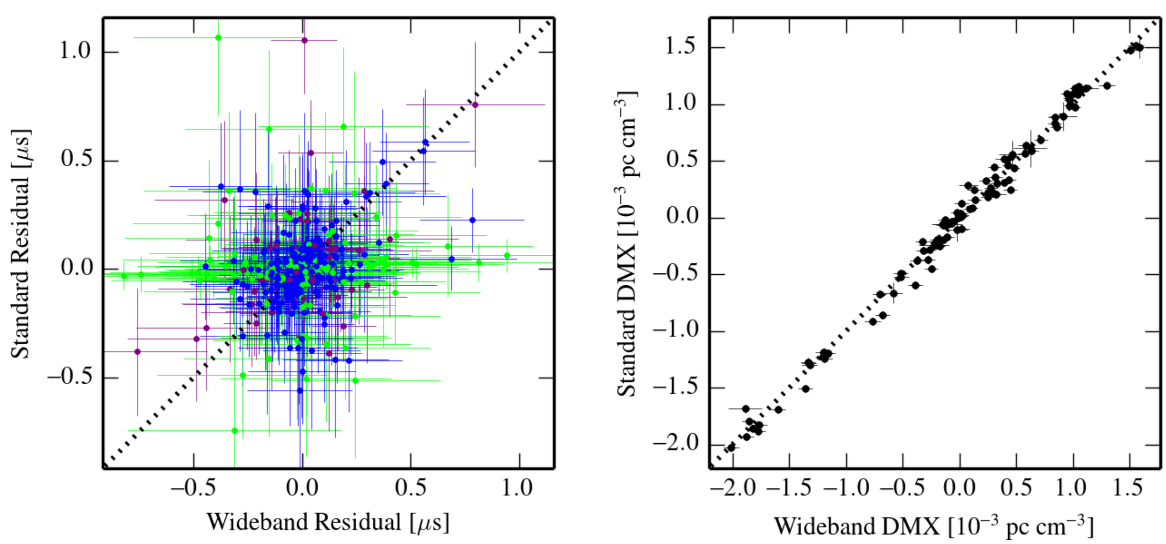

Fig. 4.98 – Comparison subplots for B1937+21.

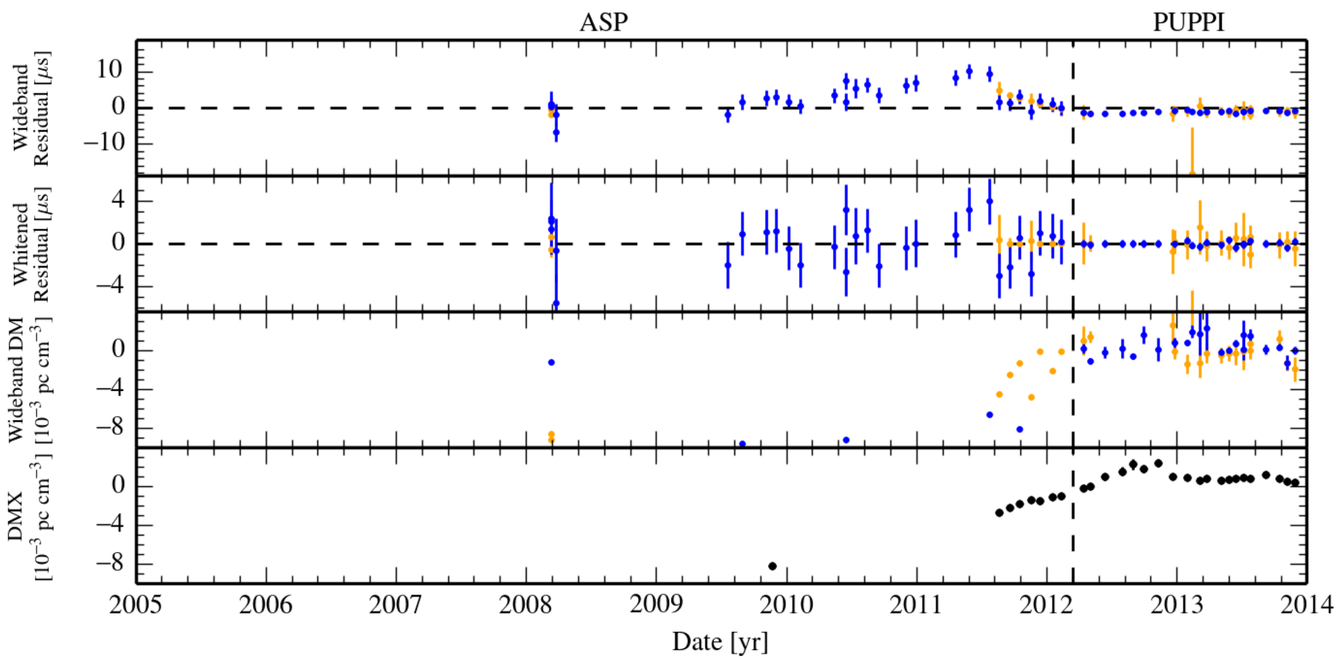

Fig. 4.99 – Timing summary panels for J1944+0907.

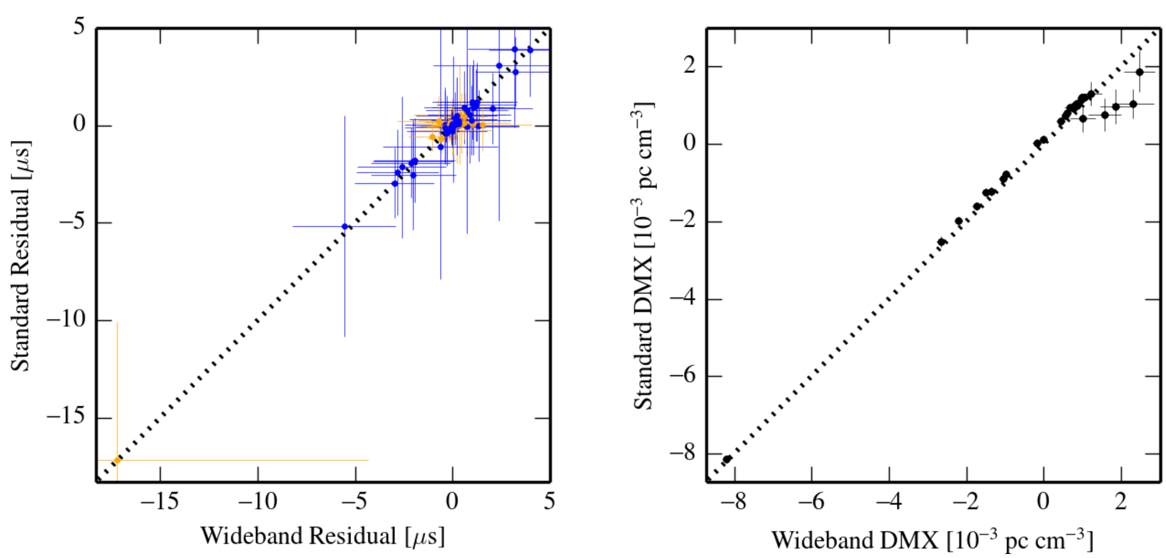

Fig. 4.100 – Comparison subplots for J1944+0907.

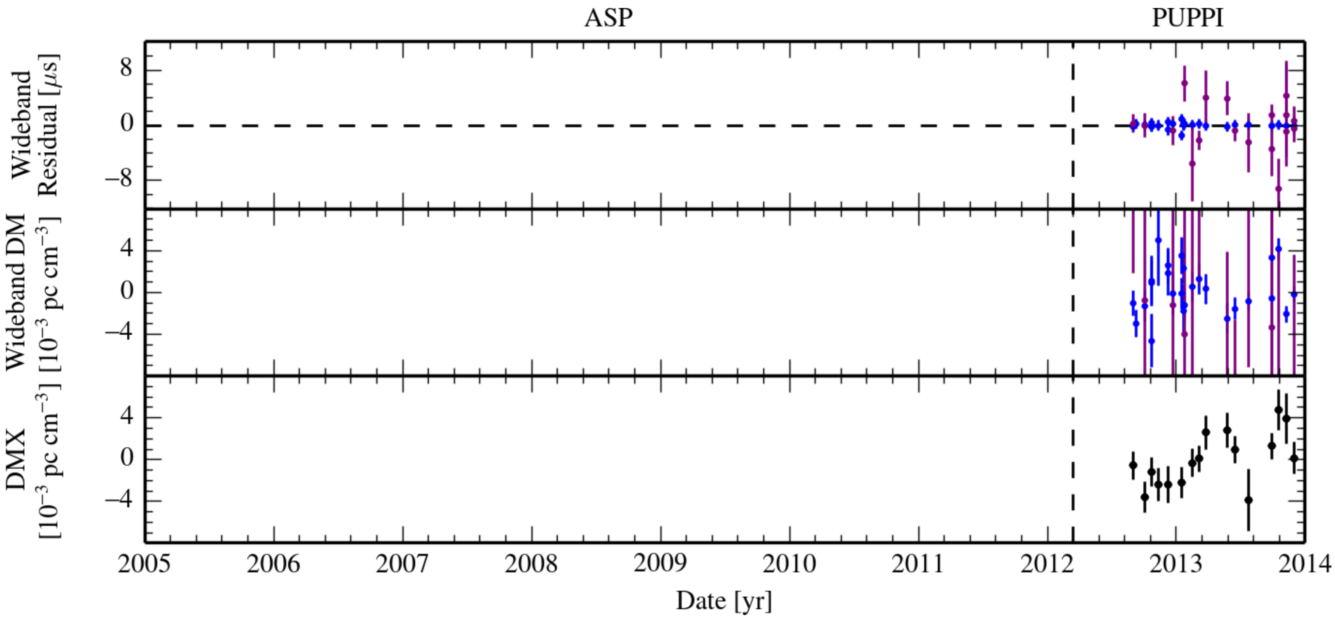

Fig. 4.101 – Timing summary panels for J1949+3106.

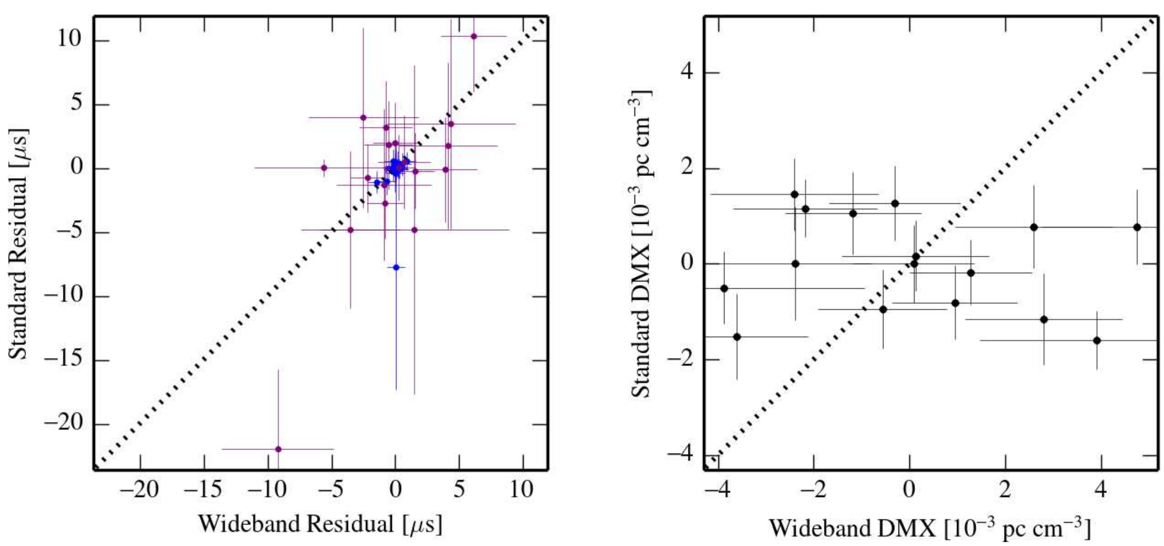

Fig. 4.102 – Comparison subplots for J1949+3106.

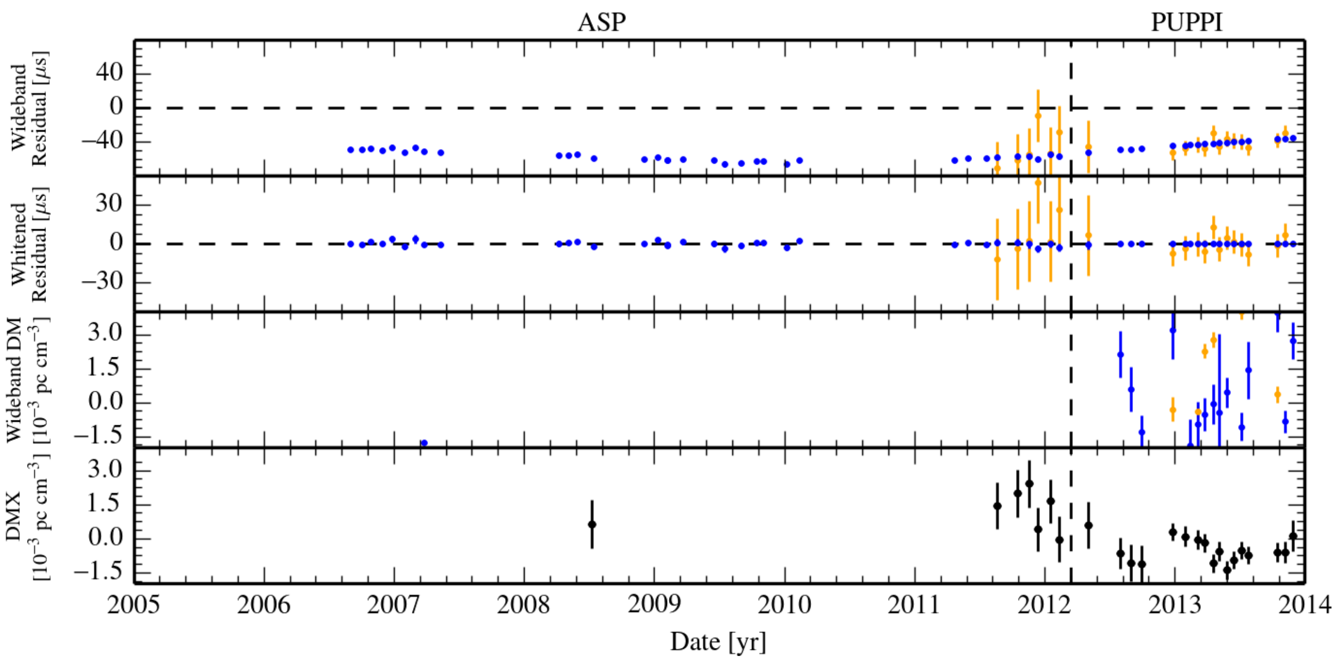

Fig. 4.103 – Timing summary panels for B1953+29.

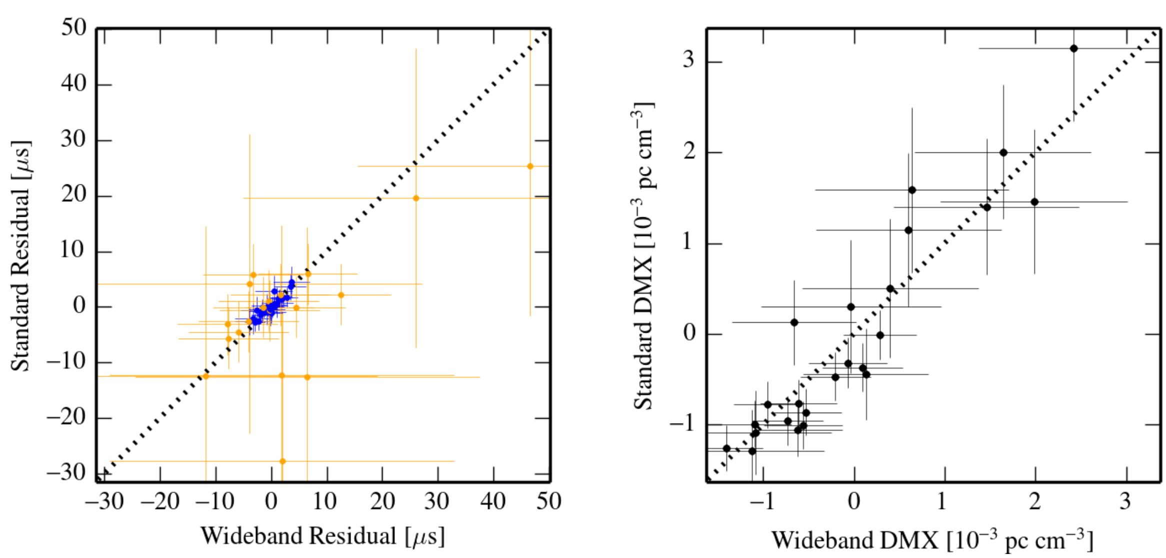

Fig. 4.104 – Comparison subplots for B1953+29.

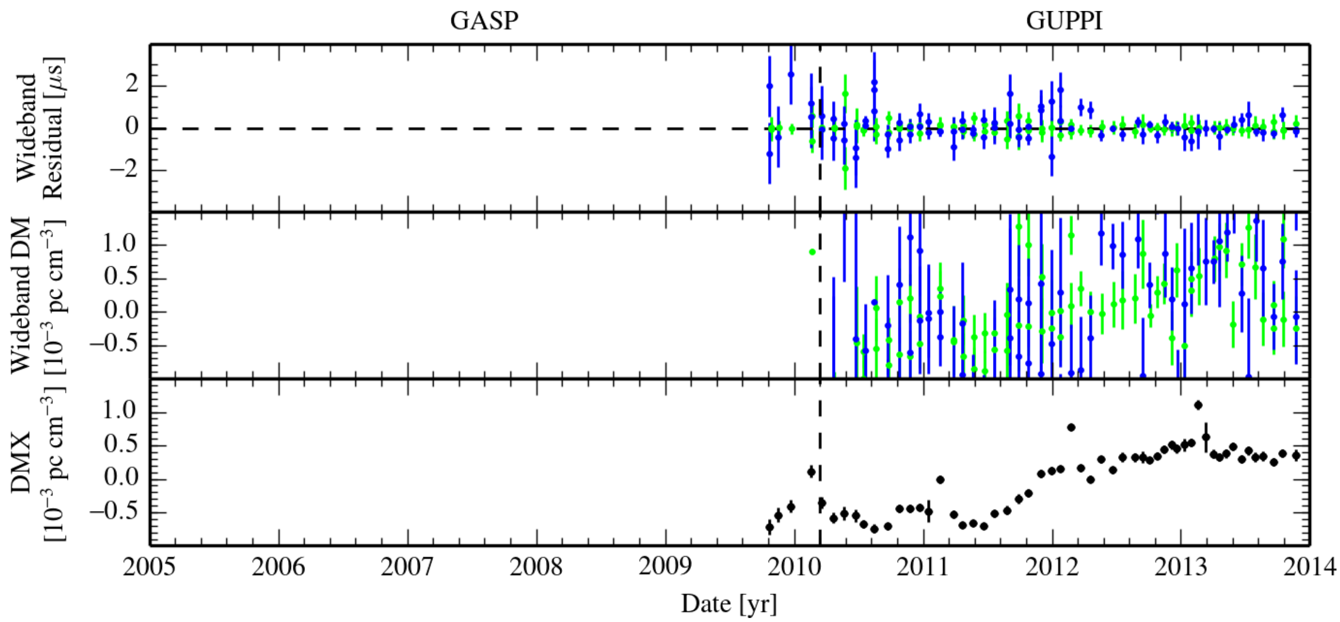

Fig. 4.105 – Timing summary panels for J2010−1323.

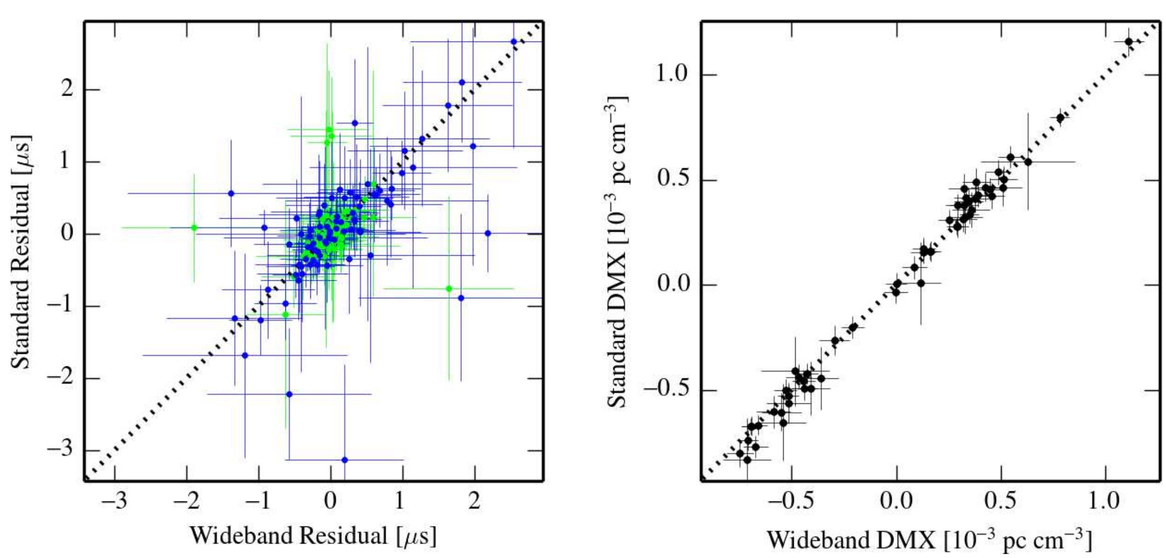

Fig. 4.106 – Comparison subplots for J2010−1323.

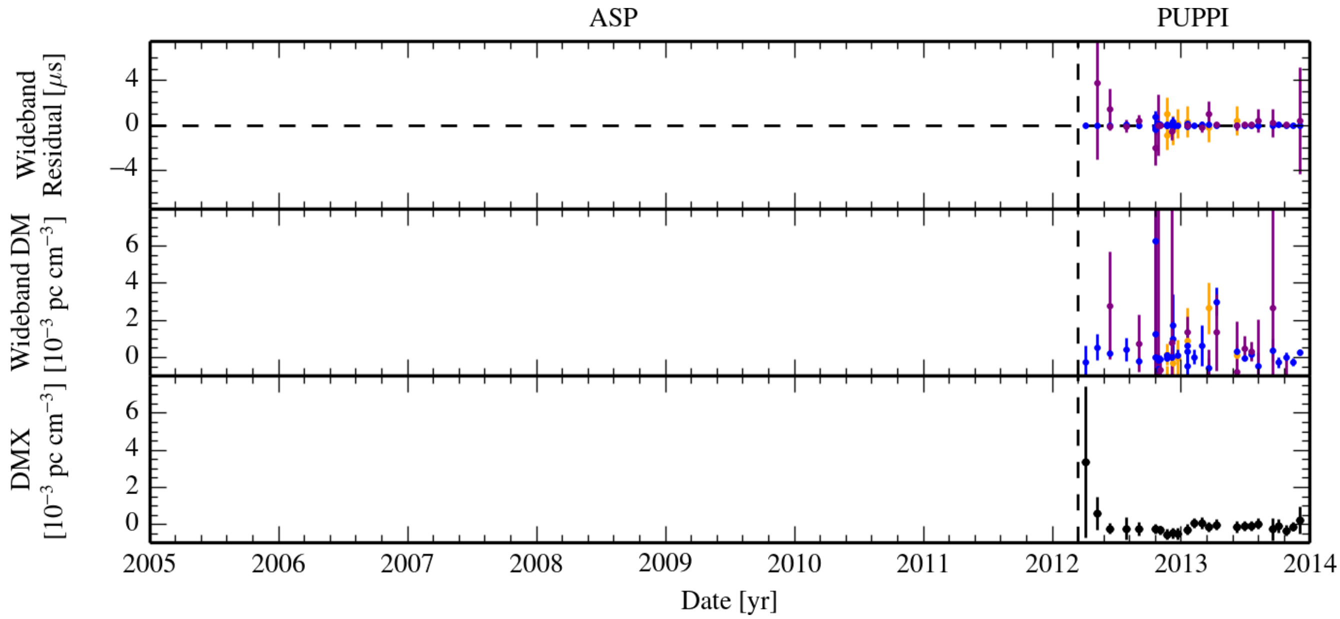

Fig. 4.107 – Timing summary panels for J2017+0603.

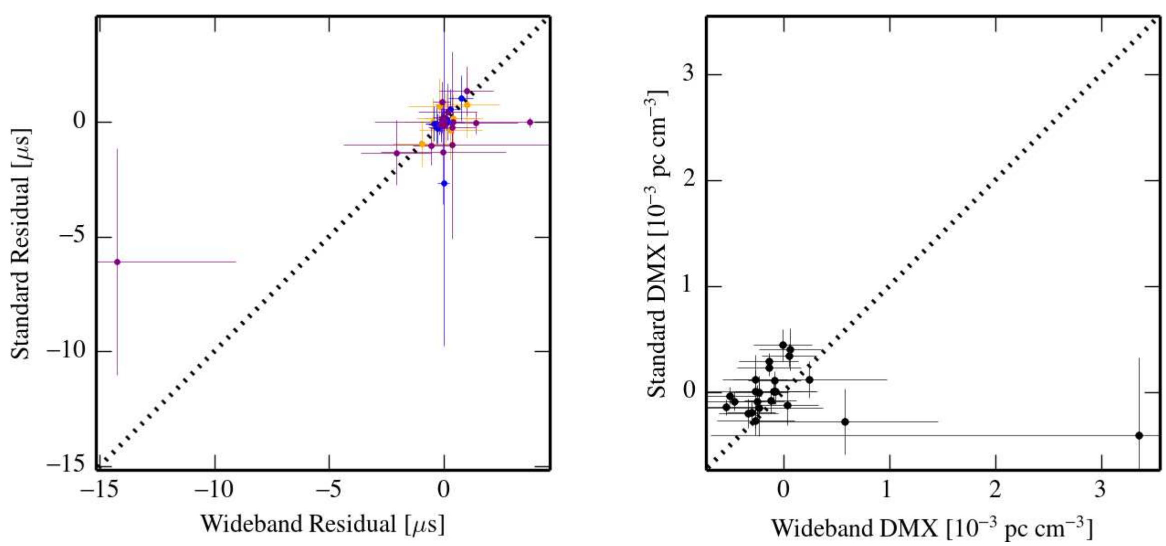

Fig. 4.108 – Comparison subplots for J2017+0603.

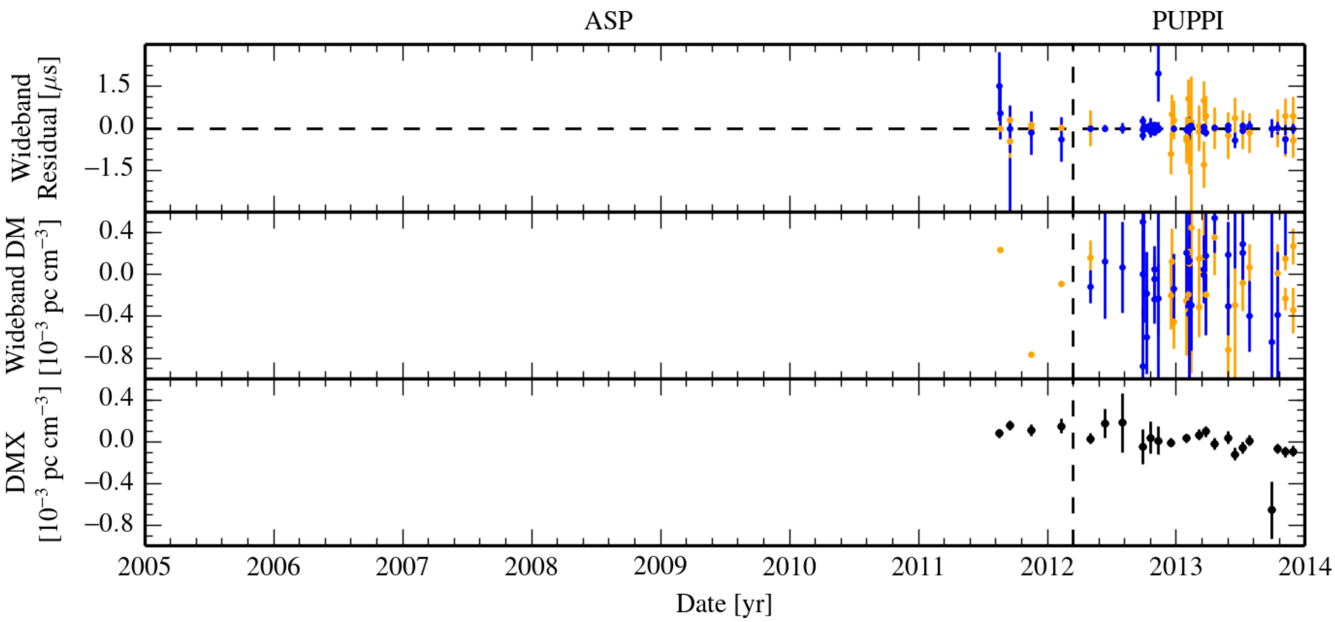

Fig. 4.109 – Timing summary panels for J2043+1711.

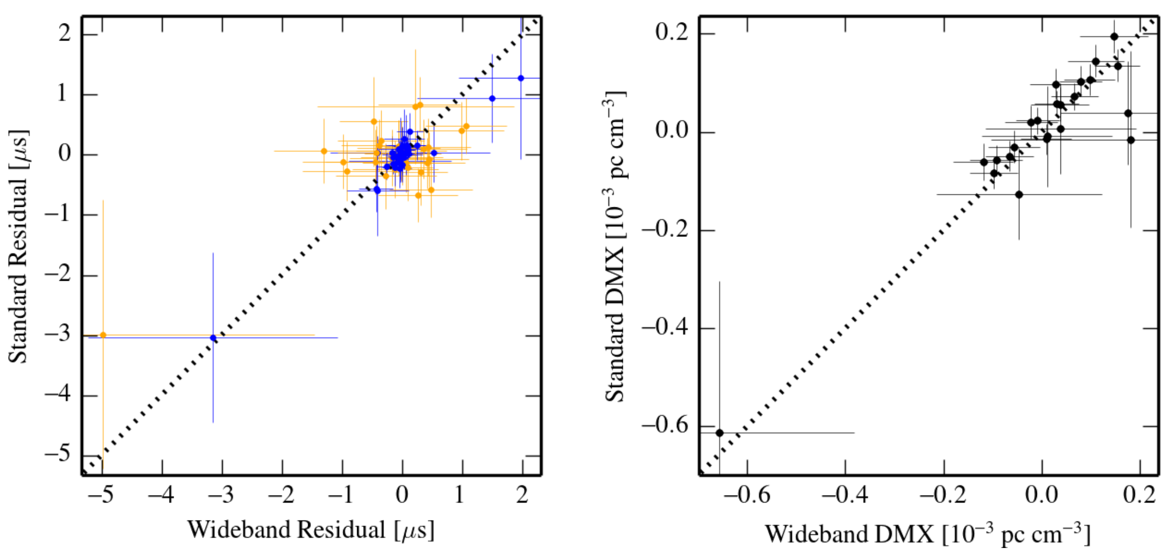

Fig. 4.110 – Comparison subplots for J2043+1711.

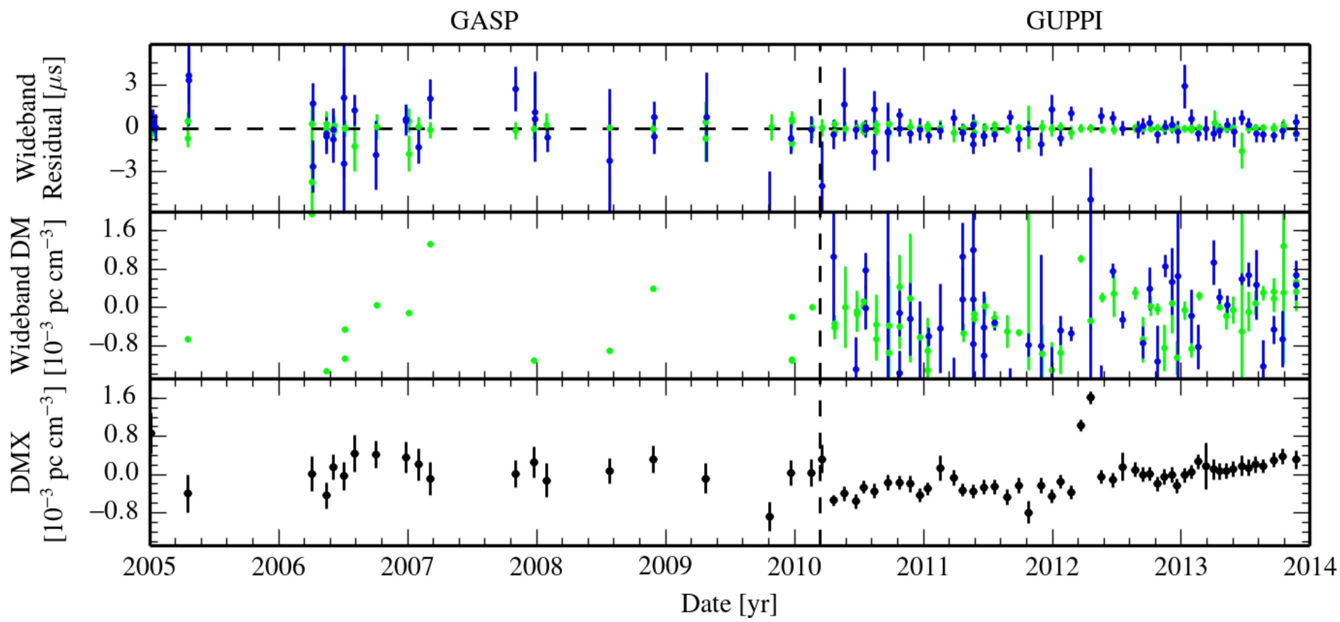

Fig. 4.111 – Timing summary panels for J2145−0750.

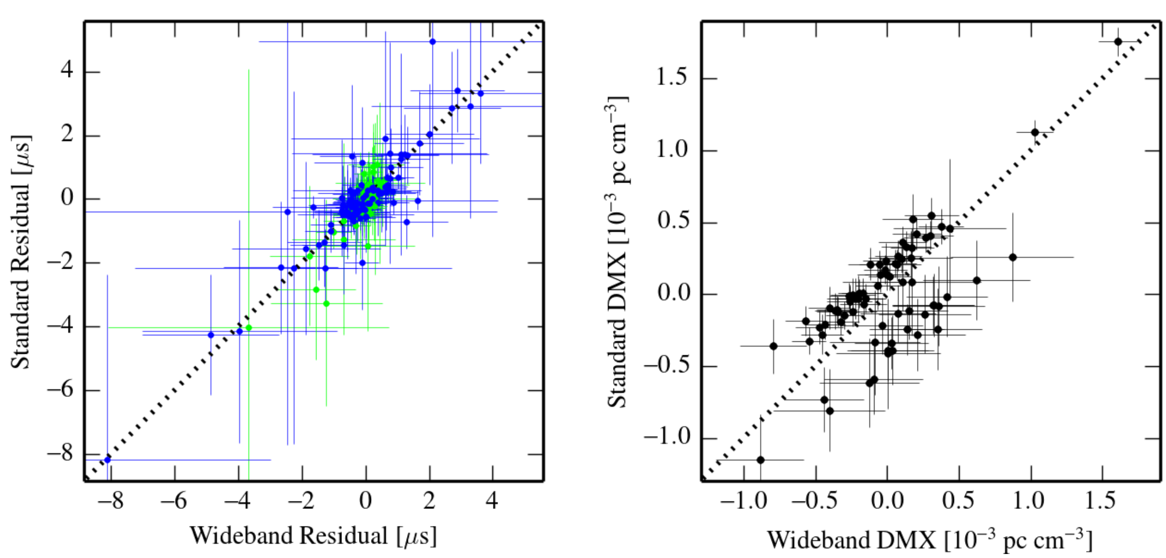

Fig. 4.112 – Comparison subplots for J2145−0750.

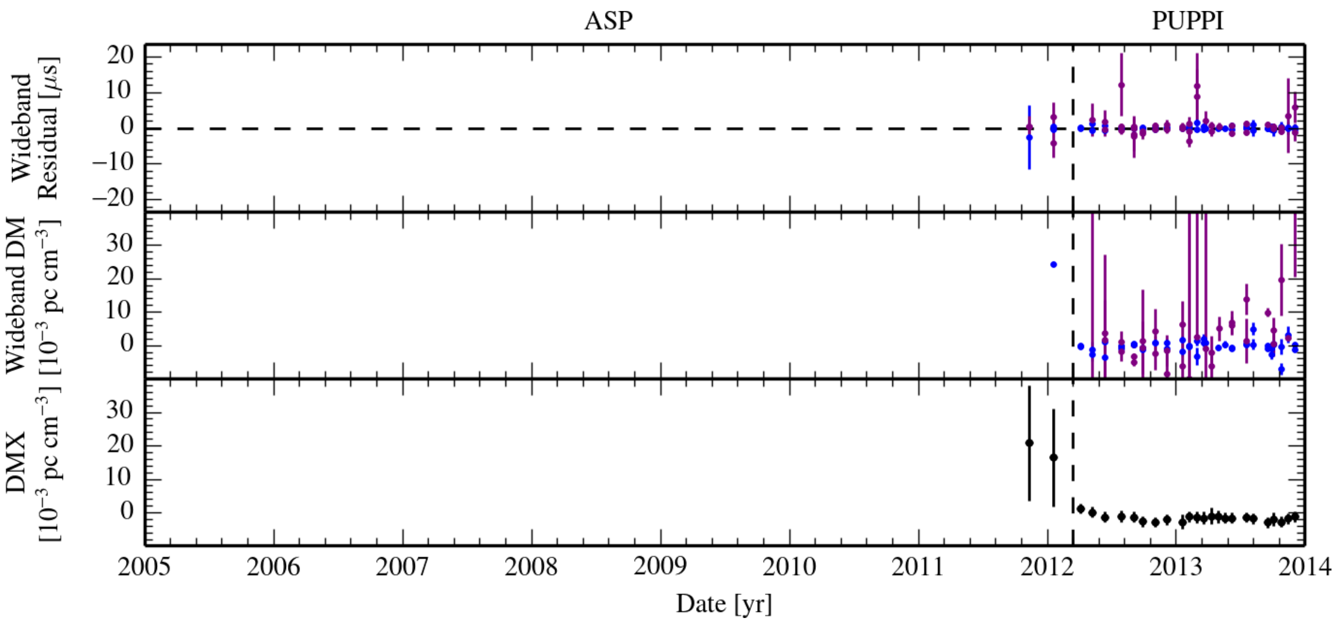

Fig. 4.113 – Timing summary panels for J2214+3000.

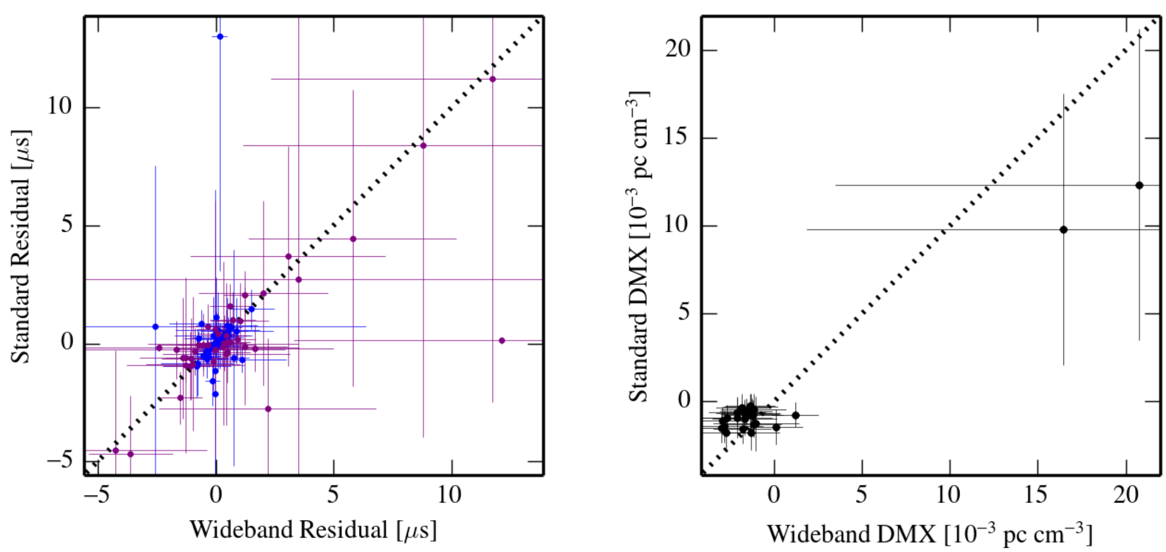

Fig. 4.114 – Comparison subplots for J2214+3000.

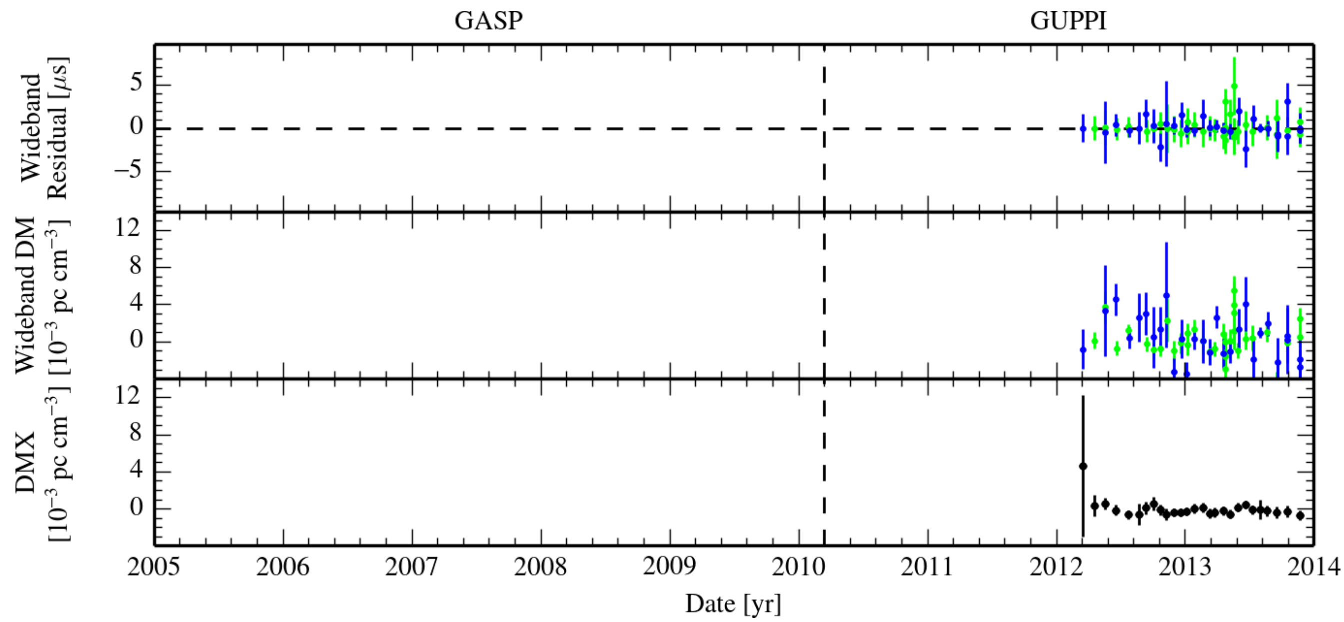

Fig. 4.115 – Timing summary panels for J2302+4442.

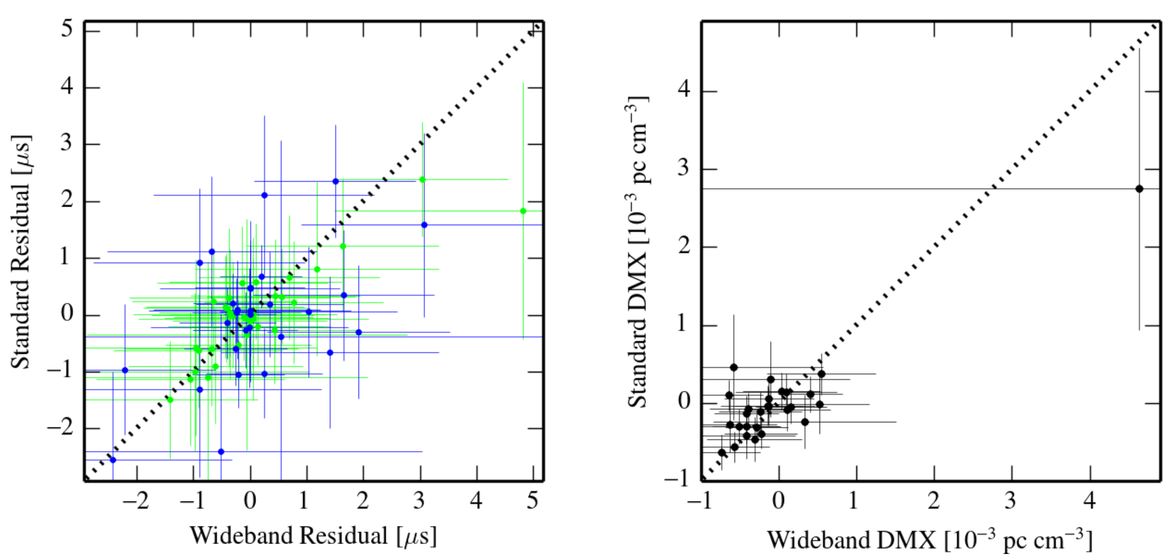

Fig. 4.116 – Comparison subplots for J2302+4442.

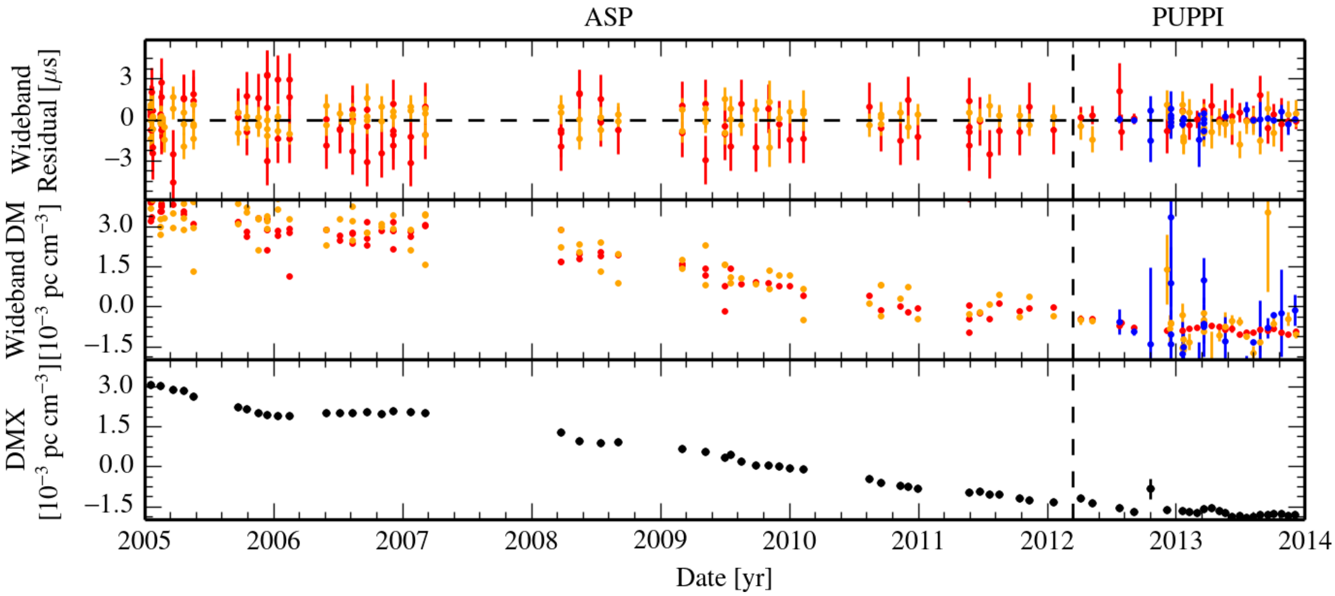

Fig. 4.117 – Timing summary panels for J2317+1439.

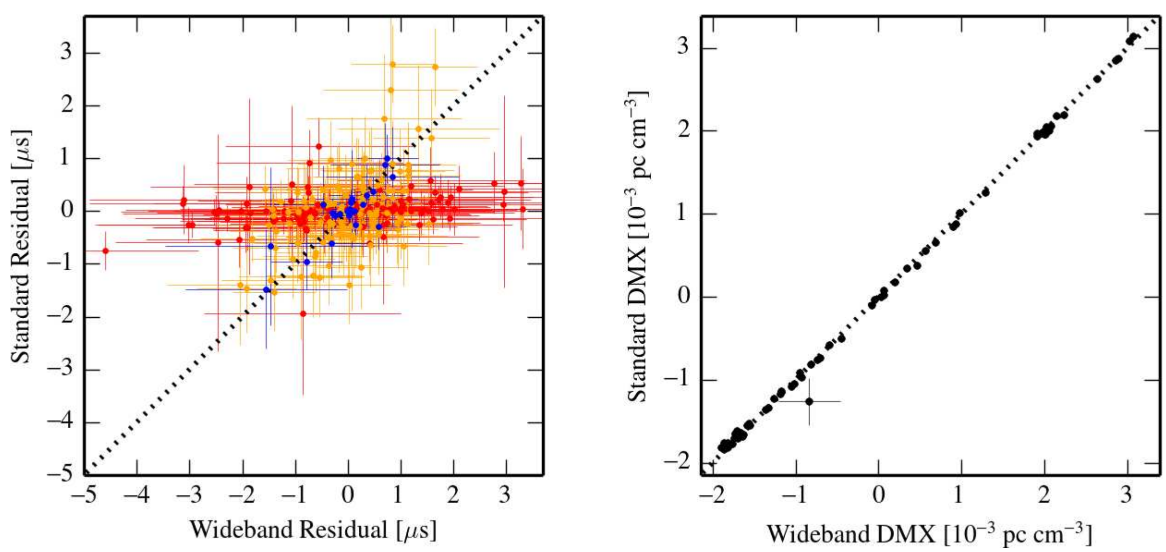

Fig. 4.118 – Comparison subplots for J2317+1439.

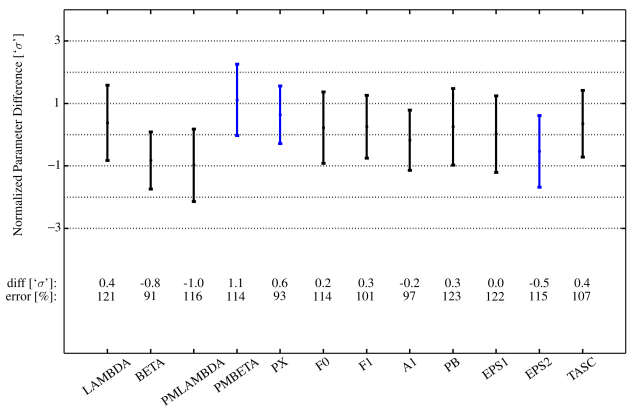

Fig. 4.119 – J0023+0923

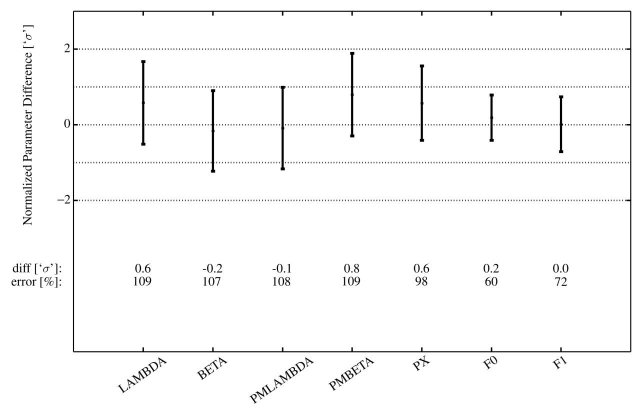

Fig. 4.120 – J0030+0451

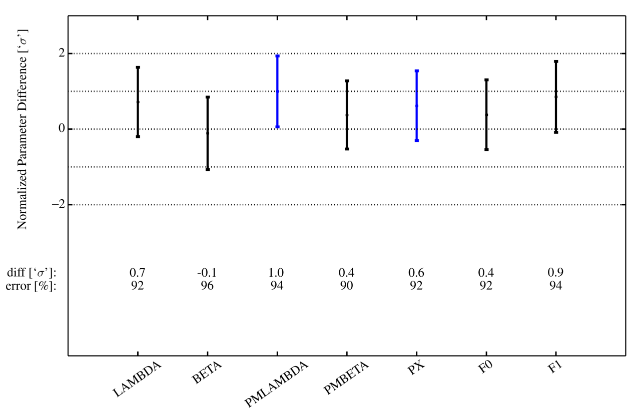

Fig. 4.121 – J0340+4130

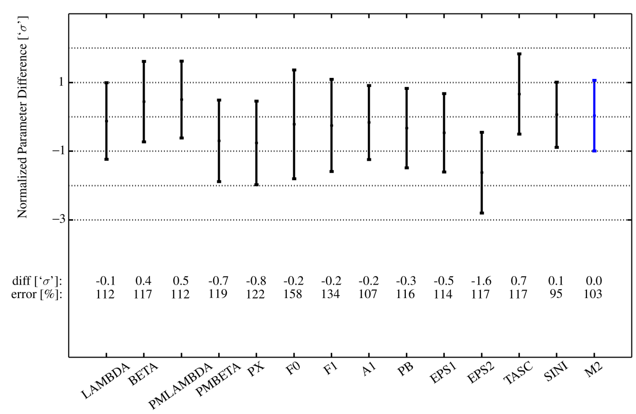

Fig. 4.122 – J0613−0200

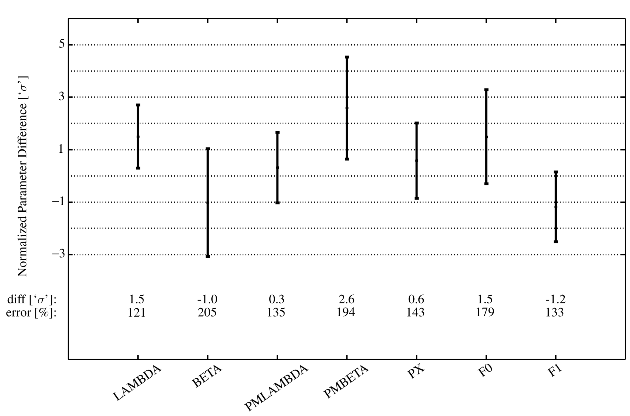

Fig. 4.123 – J0645+5158

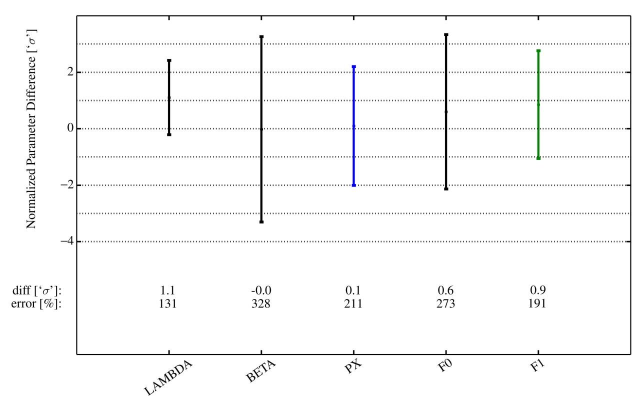

Fig. 4.124 – J0931−1902

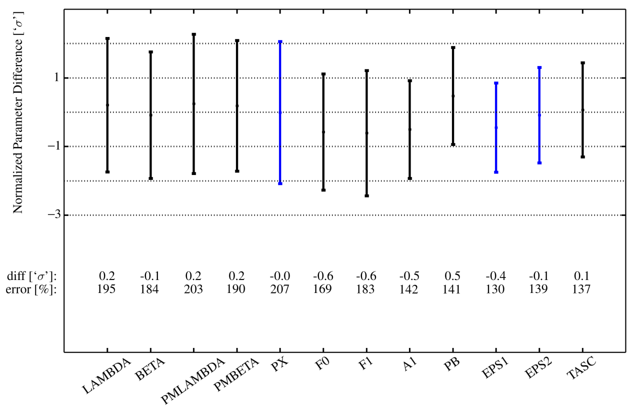

Fig. 4.125 – J1012+5307

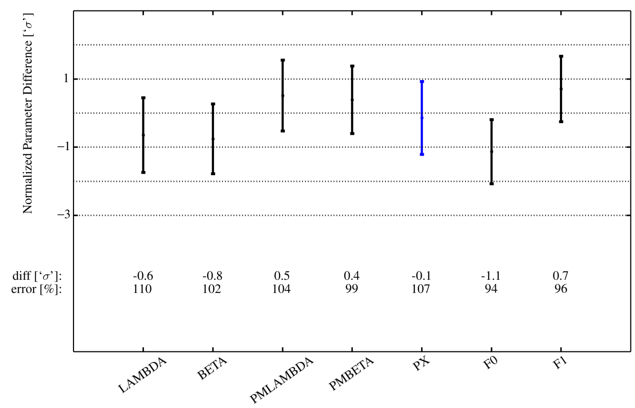

Fig. 4.126 – J1024−0719

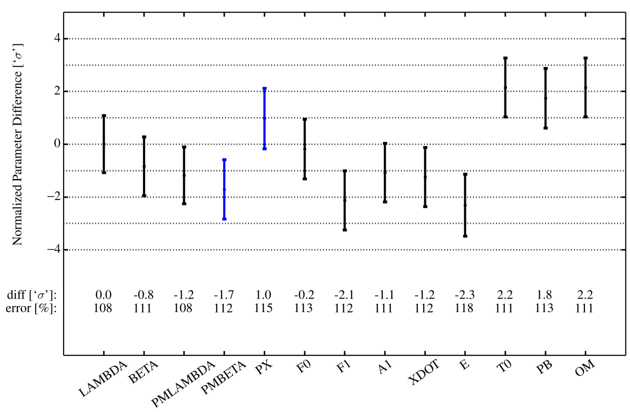

Fig. 4.127 – J1455−3330

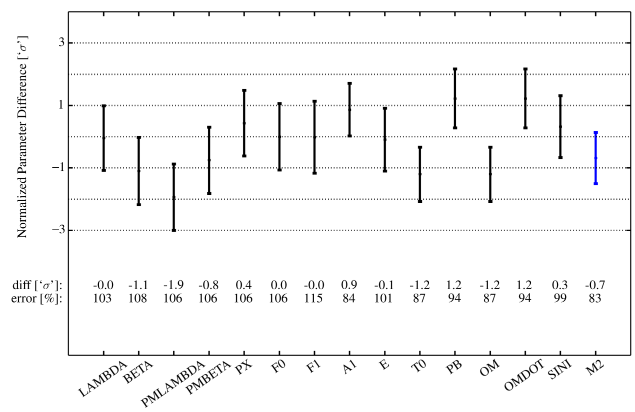

Fig. 4.128 – J1600−3053

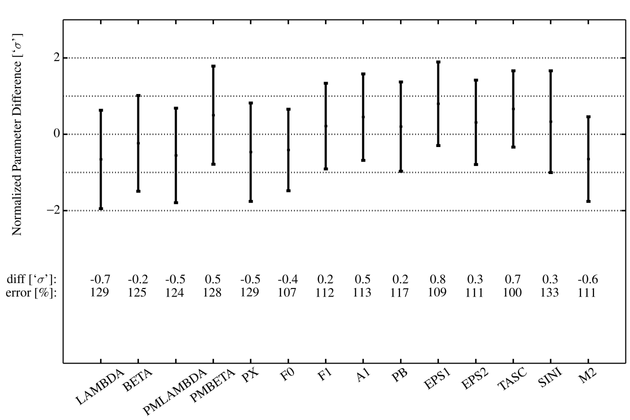

Fig. 4.129 – J1614−2230

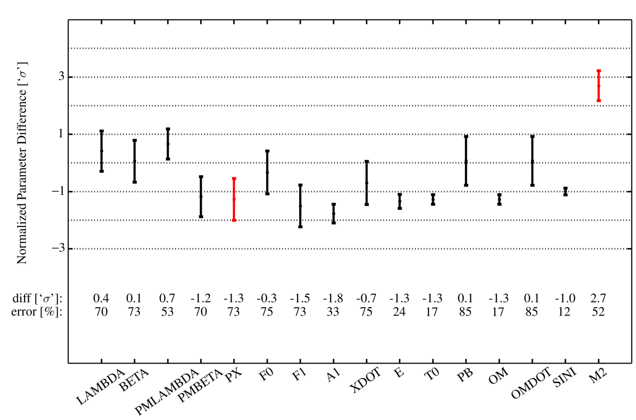

Fig. 4.130 – J1640+2224

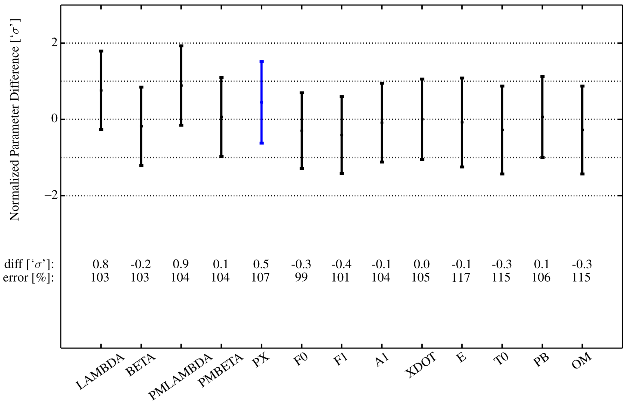

Fig. 4.131 – J1643−1224

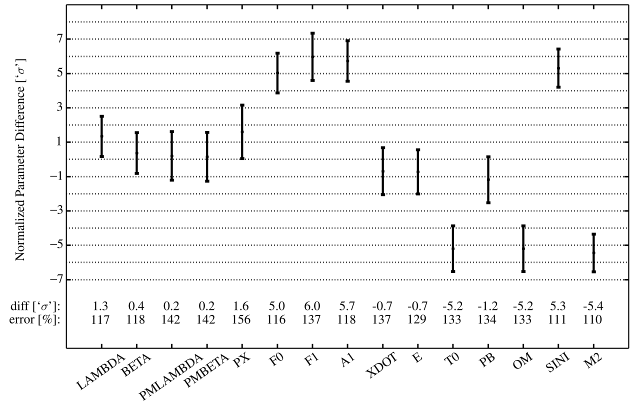

Fig. 4.132 – J1713+0747

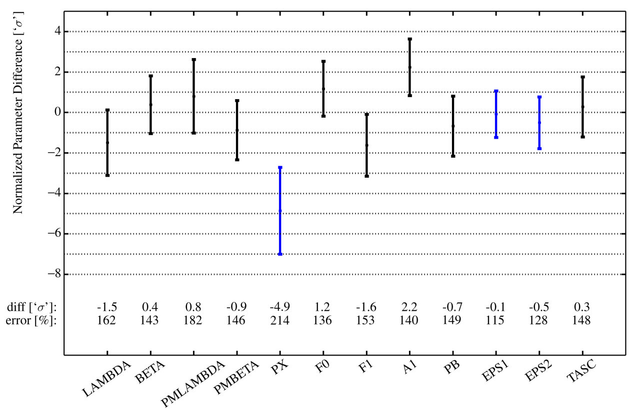

Fig. 4.133 – J1738+0333

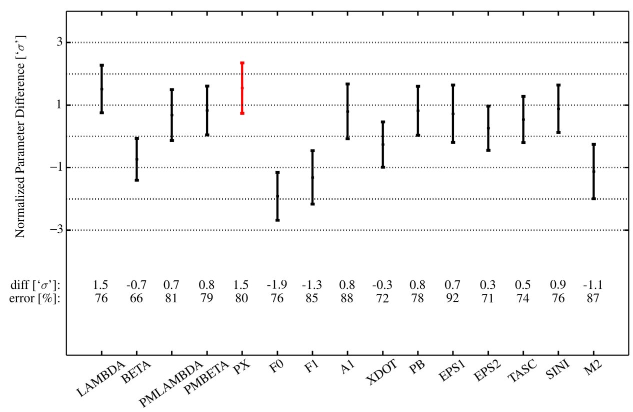

Fig. 4.134 – J1741+1351

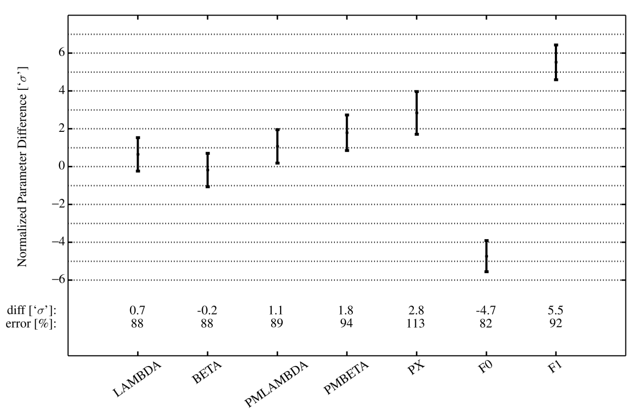

Fig. 4.135 – J1744−1134

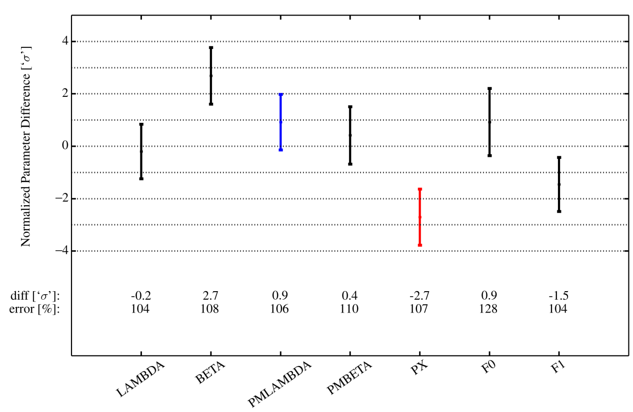

Fig. 4.136 – J1747−4036

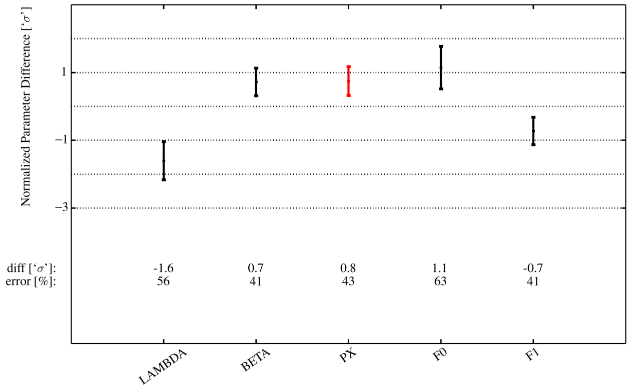

Fig. 4.137 – J1832−0836

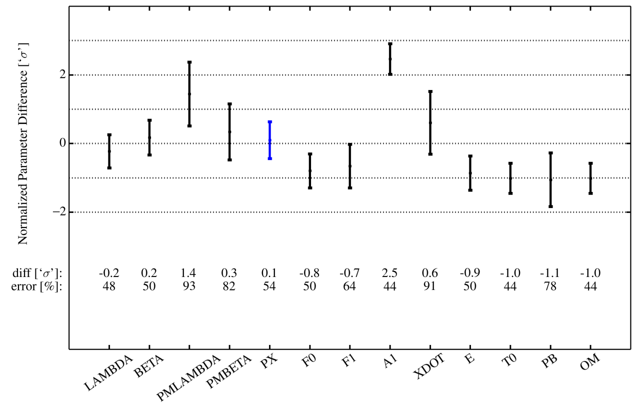

Fig. 4.138 – J1853+1303

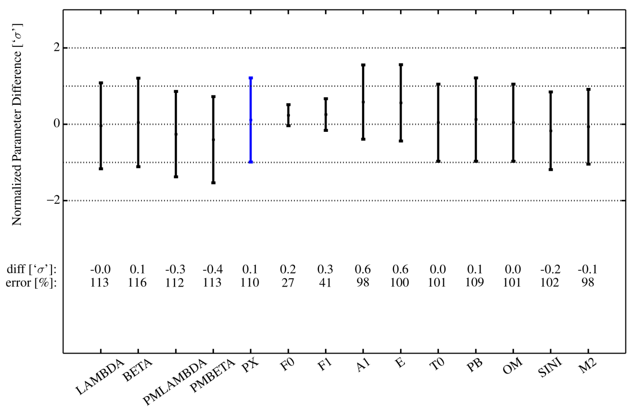

Fig. 4.139 – B1855+09

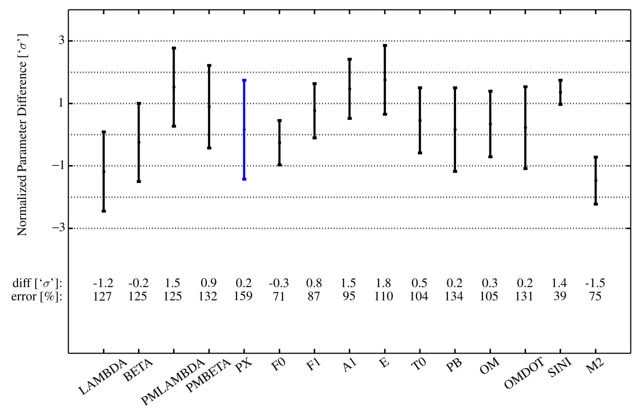

Fig. 4.140 – J1903+0327

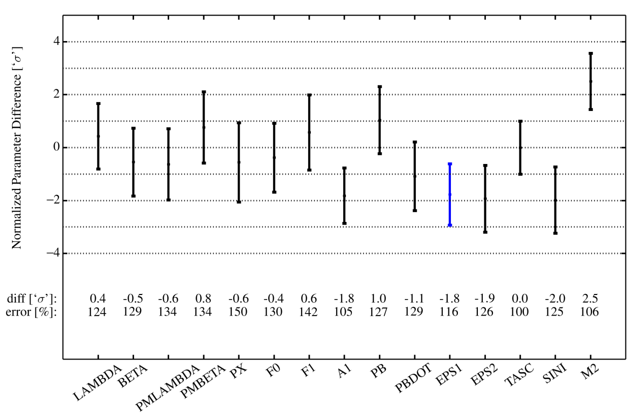

Fig. 4.141 – J1909−3744

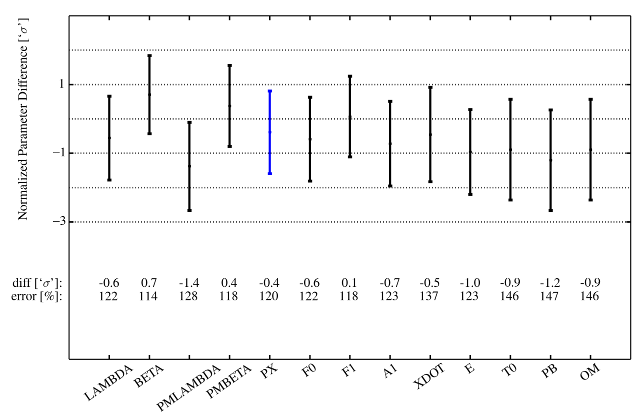

Fig. 4.142 – J1910+1256

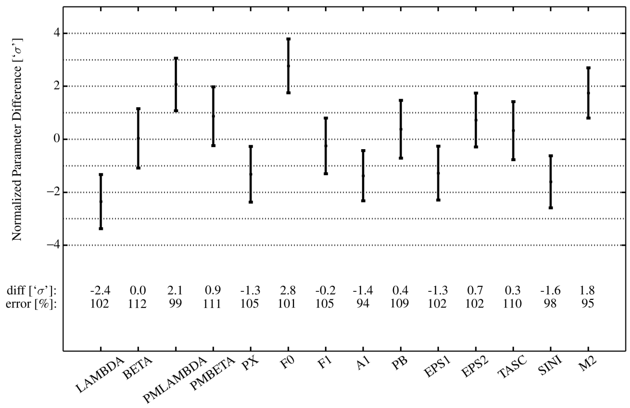

Fig. 4.143 – J1918−0642

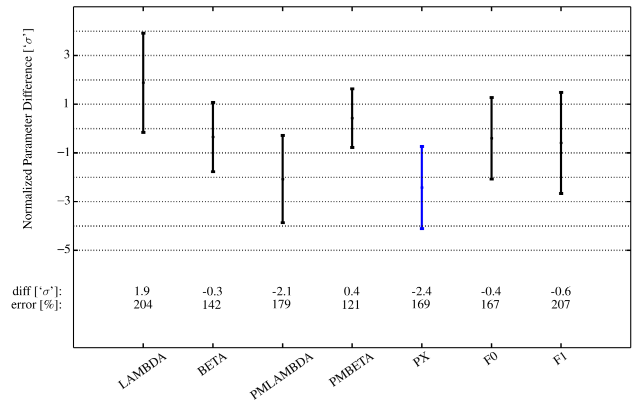

Fig. 4.144 – J1923+2515

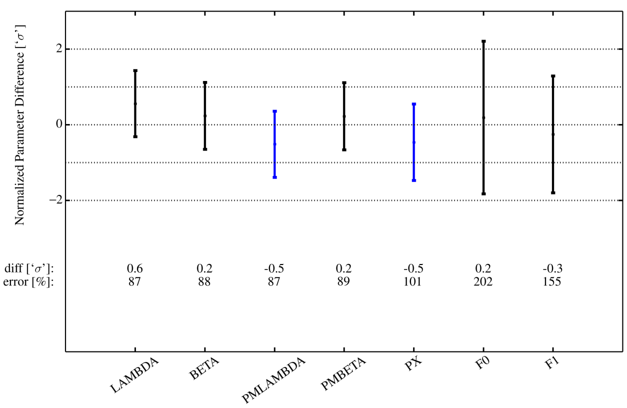

Fig. 4.145 – B1937+21

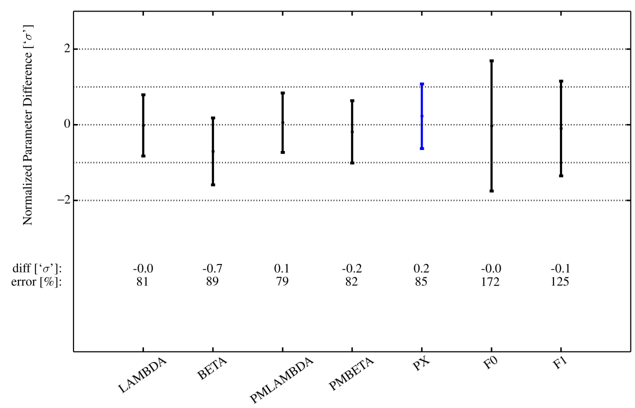

Fig. 4.146 – J1944+0907

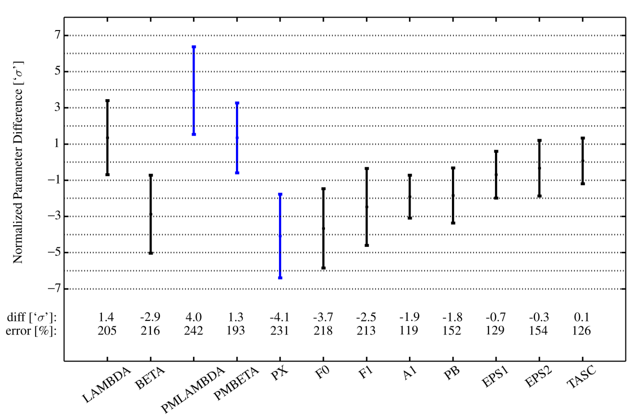

Fig. 4.147 – J1949+3106

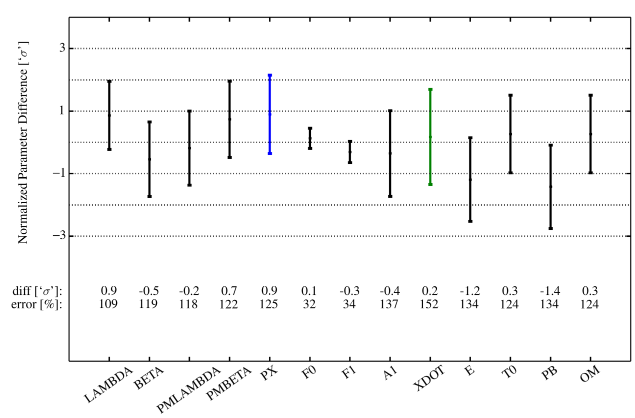

Fig. 4.148 – B1953+29

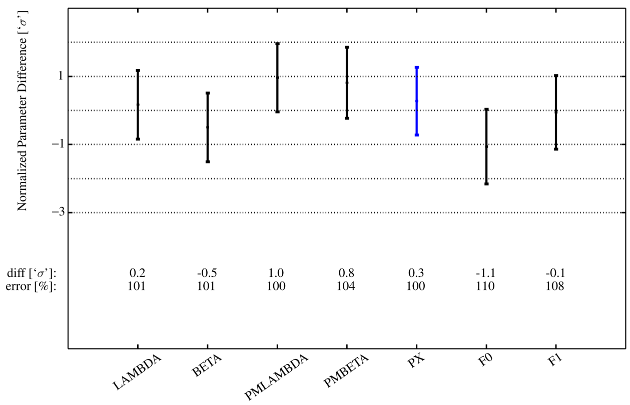

Fig. 4.149 – J2010−1323

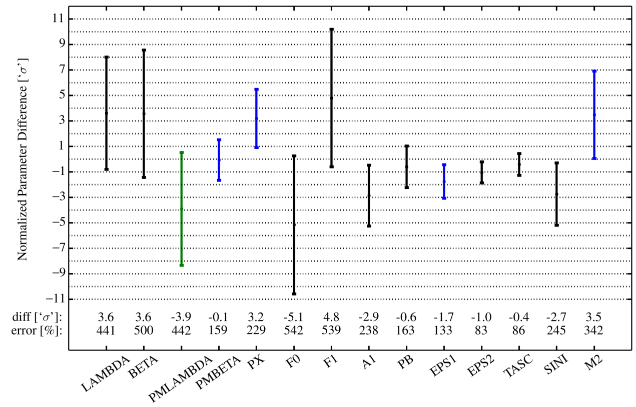

Fig. 4.150 – J2017+0603

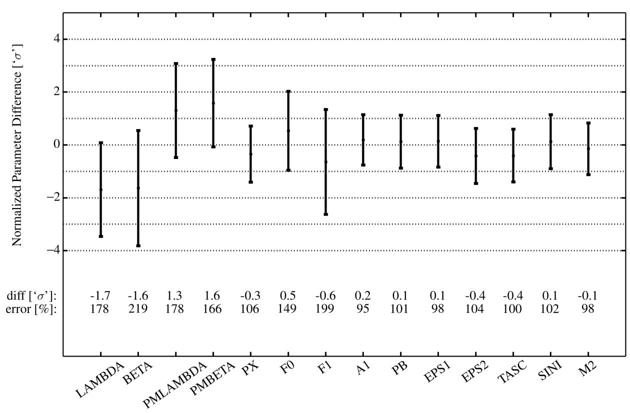

Fig. 4.151 – J2043+1711

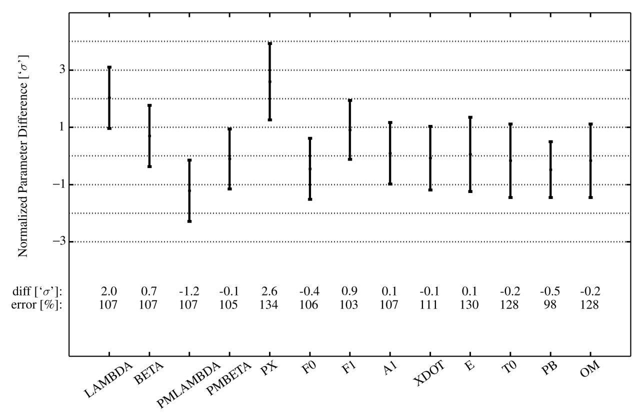

Fig. 4.152 – J2145−0750

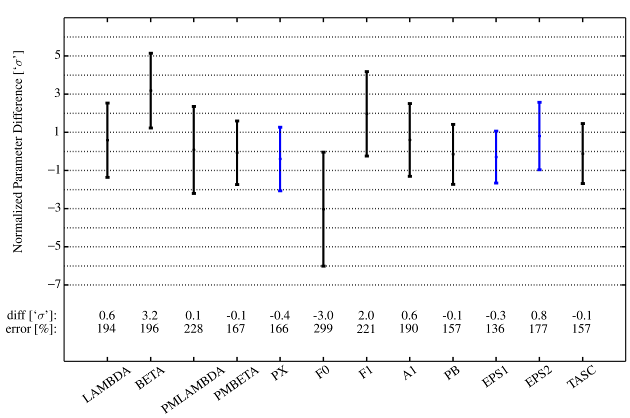

Fig. 4.153 – J2214+3000

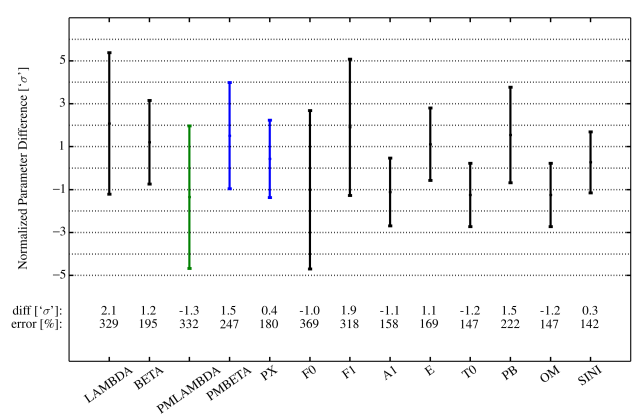

Fig. 4.154 – J2302+4442

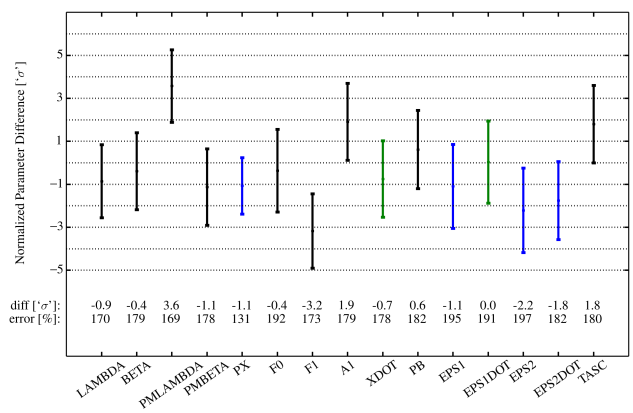

Fig. 4.155 – J2317+1439

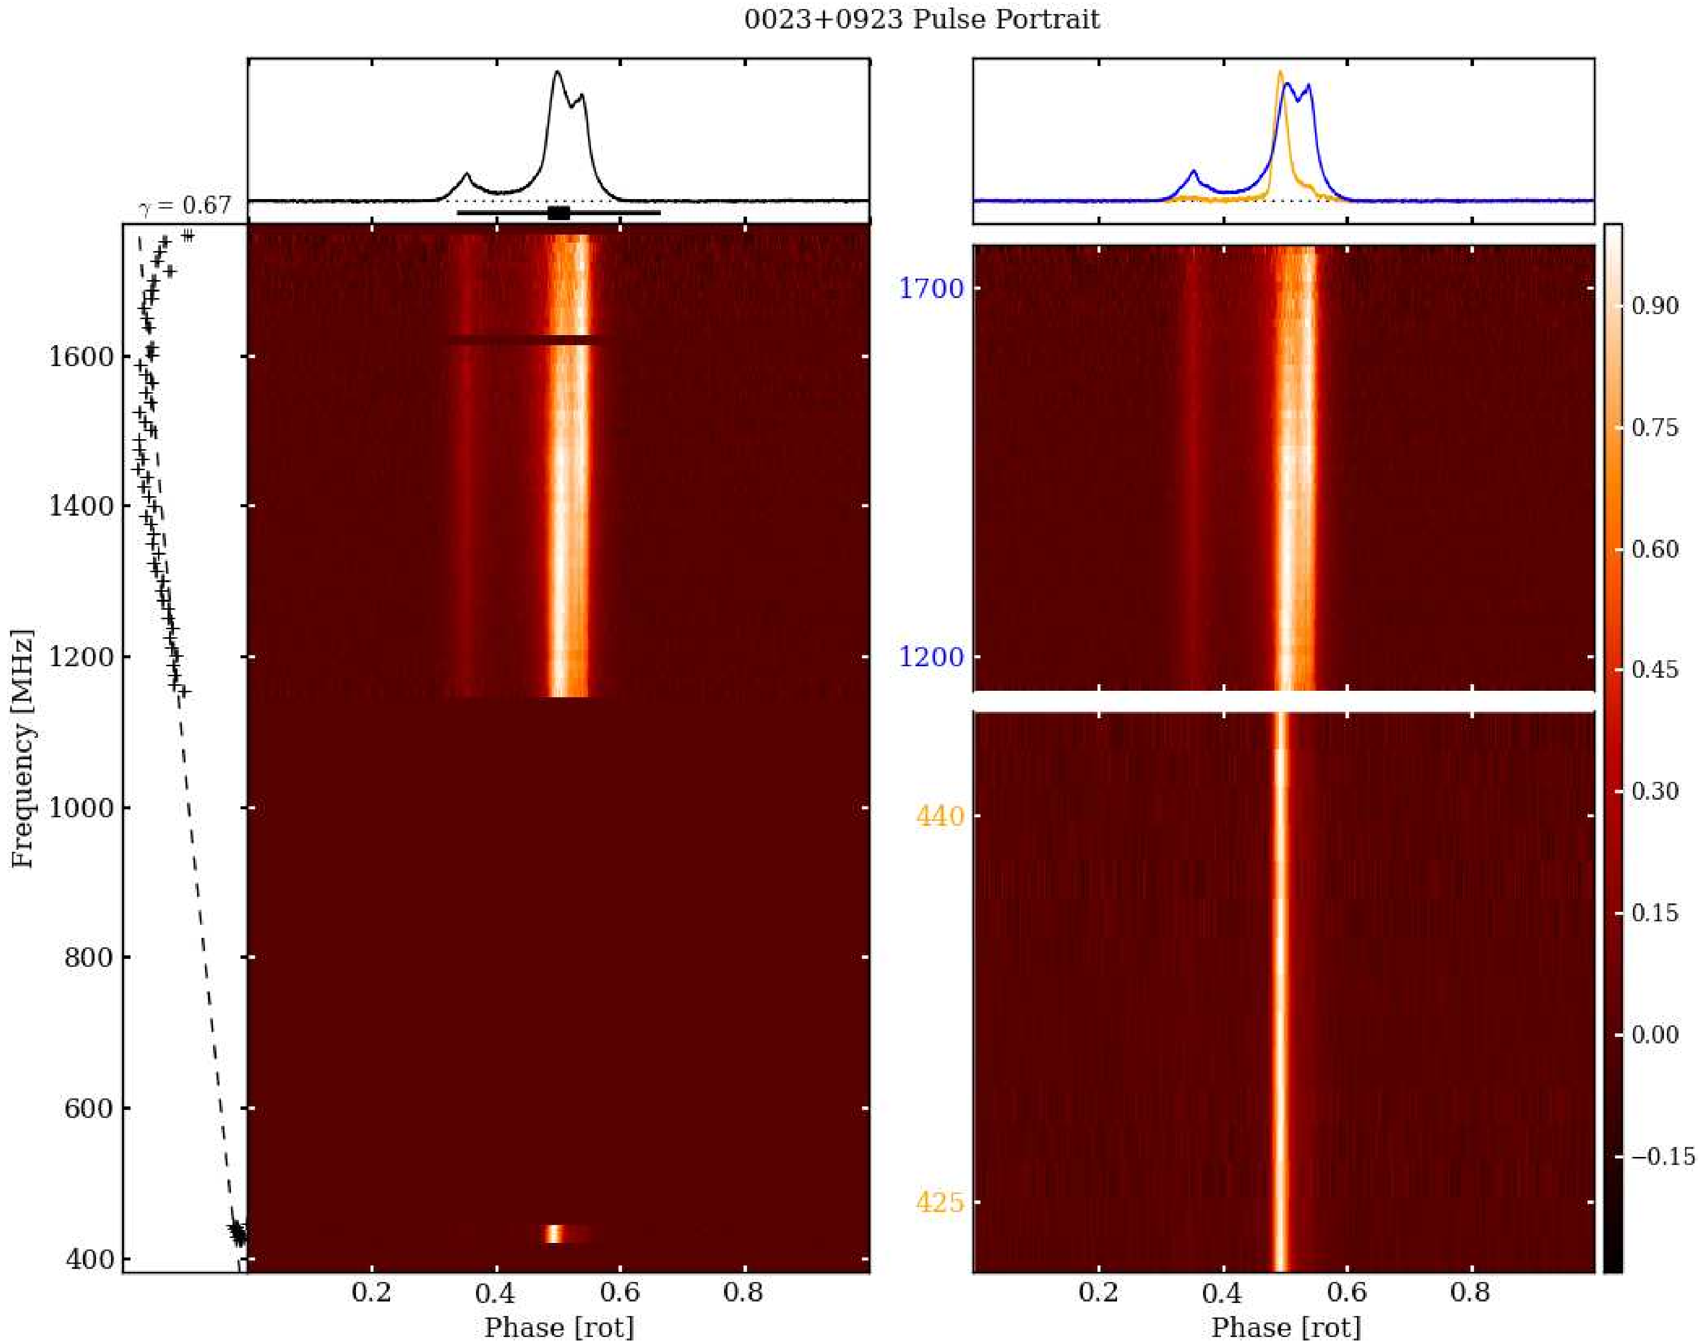

Fig. 4.156 – Concatenated portrait data and average profiles for J0023+0923.

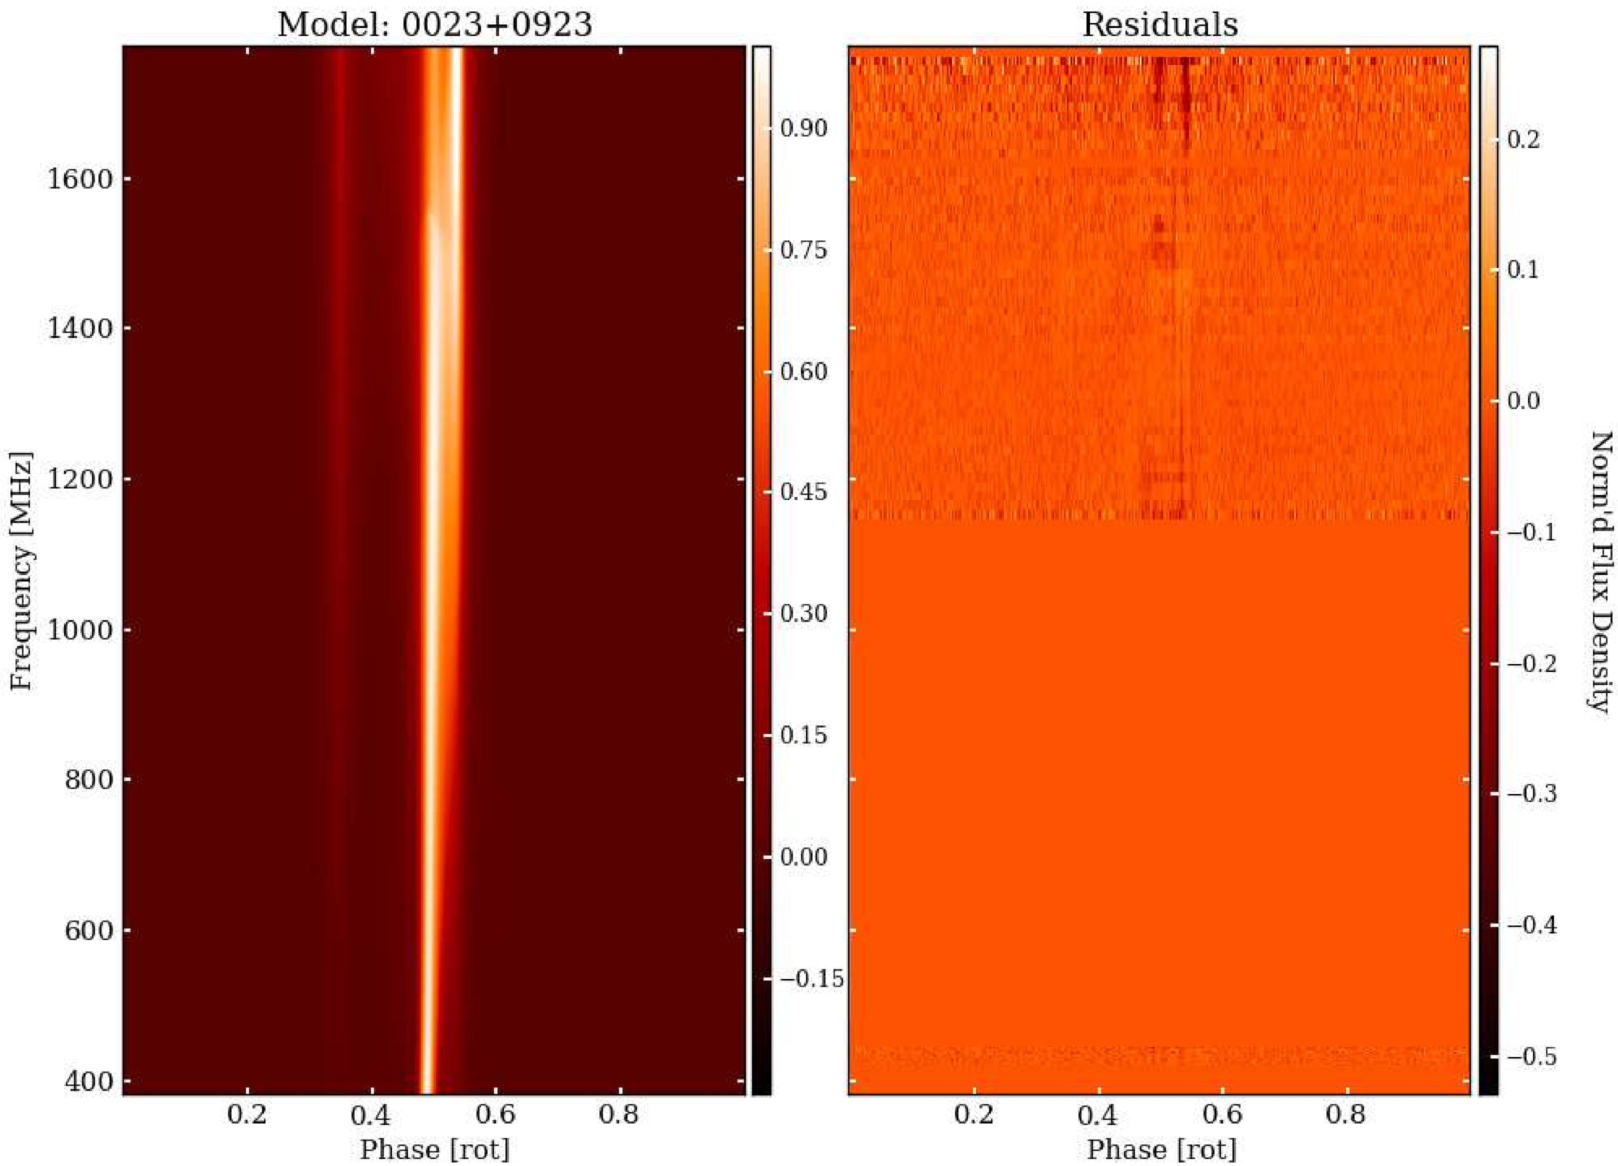

Fig. 4.157 – Fitted Gaussian component model and residuals for J0023+0923.

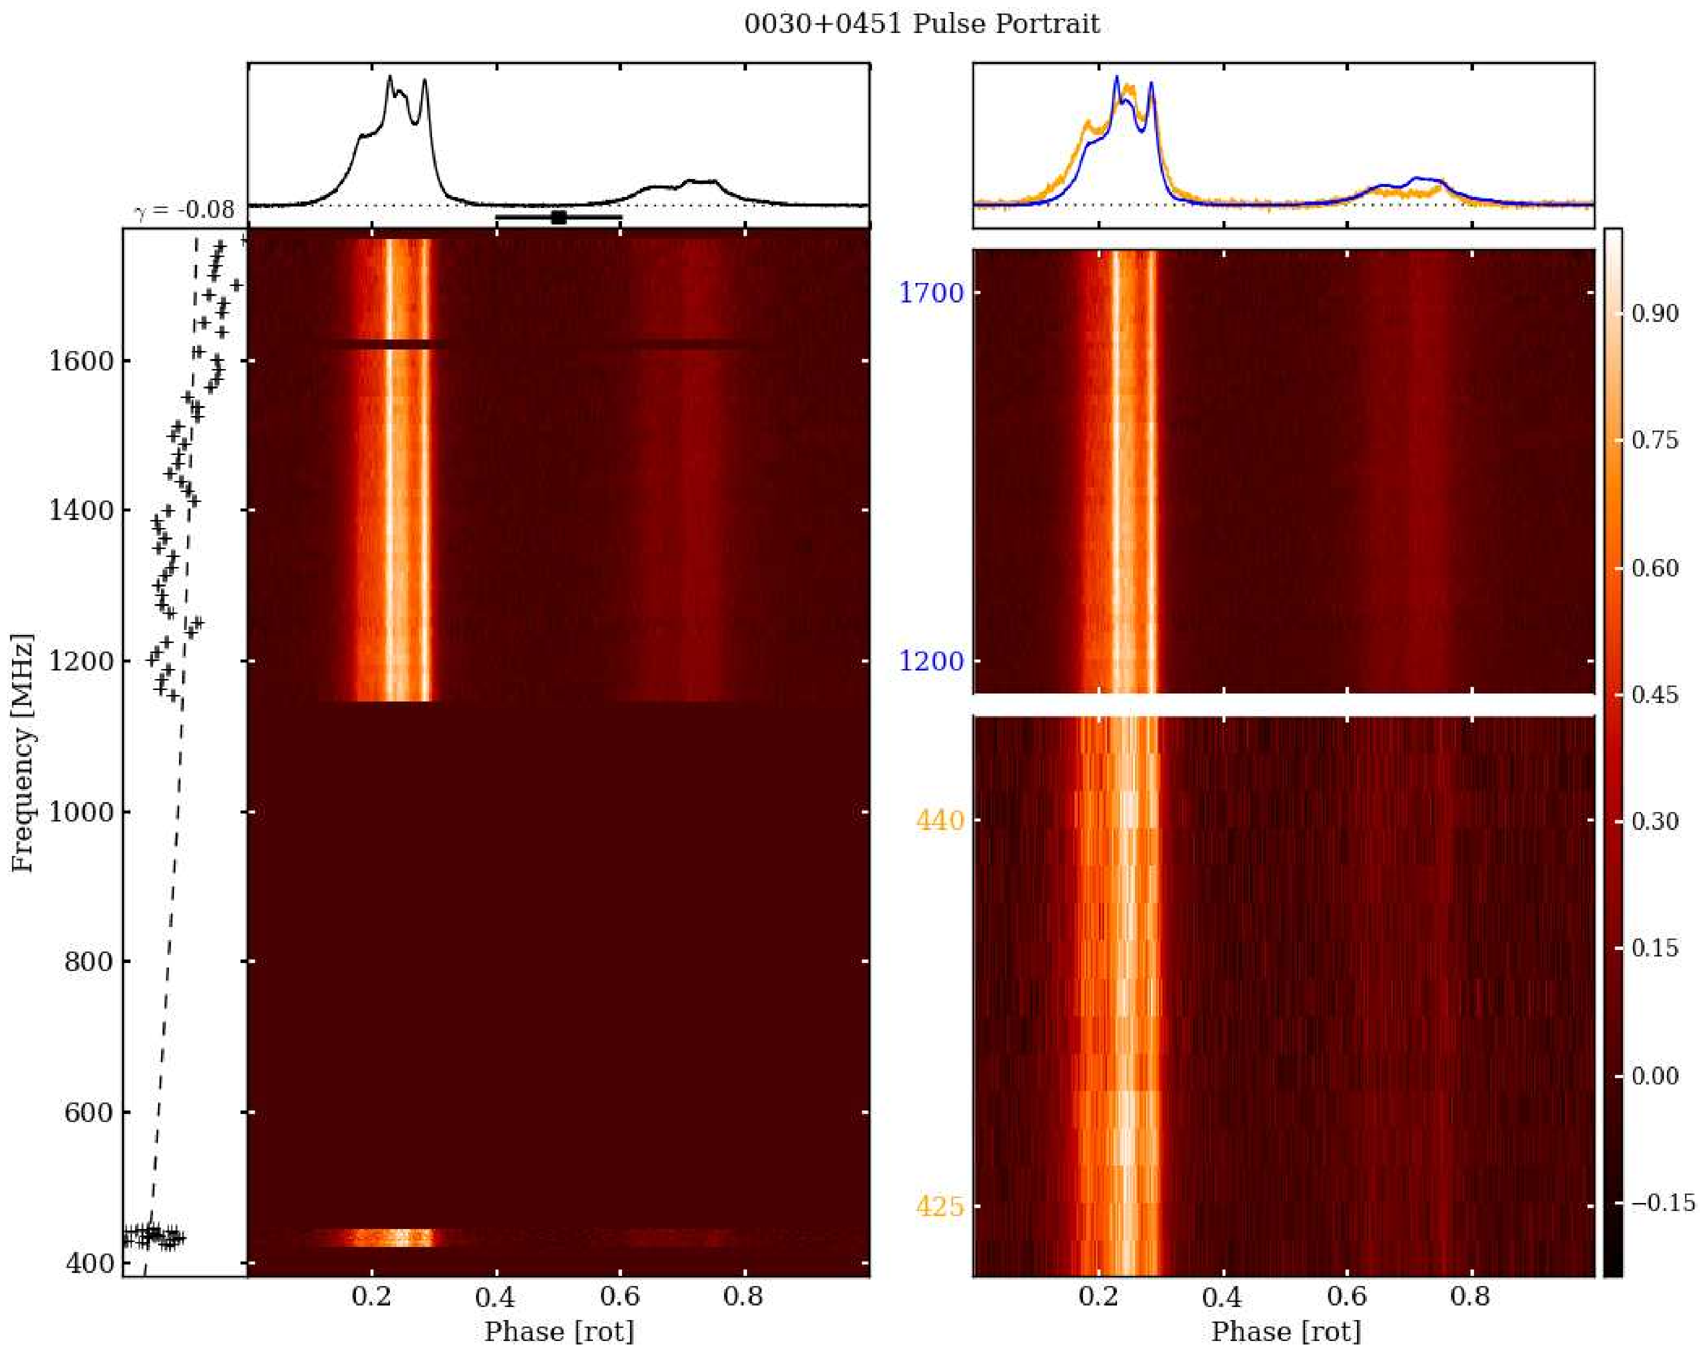

Fig. 4.158 – Concatenated portrait data and average profiles for J0030+0451.

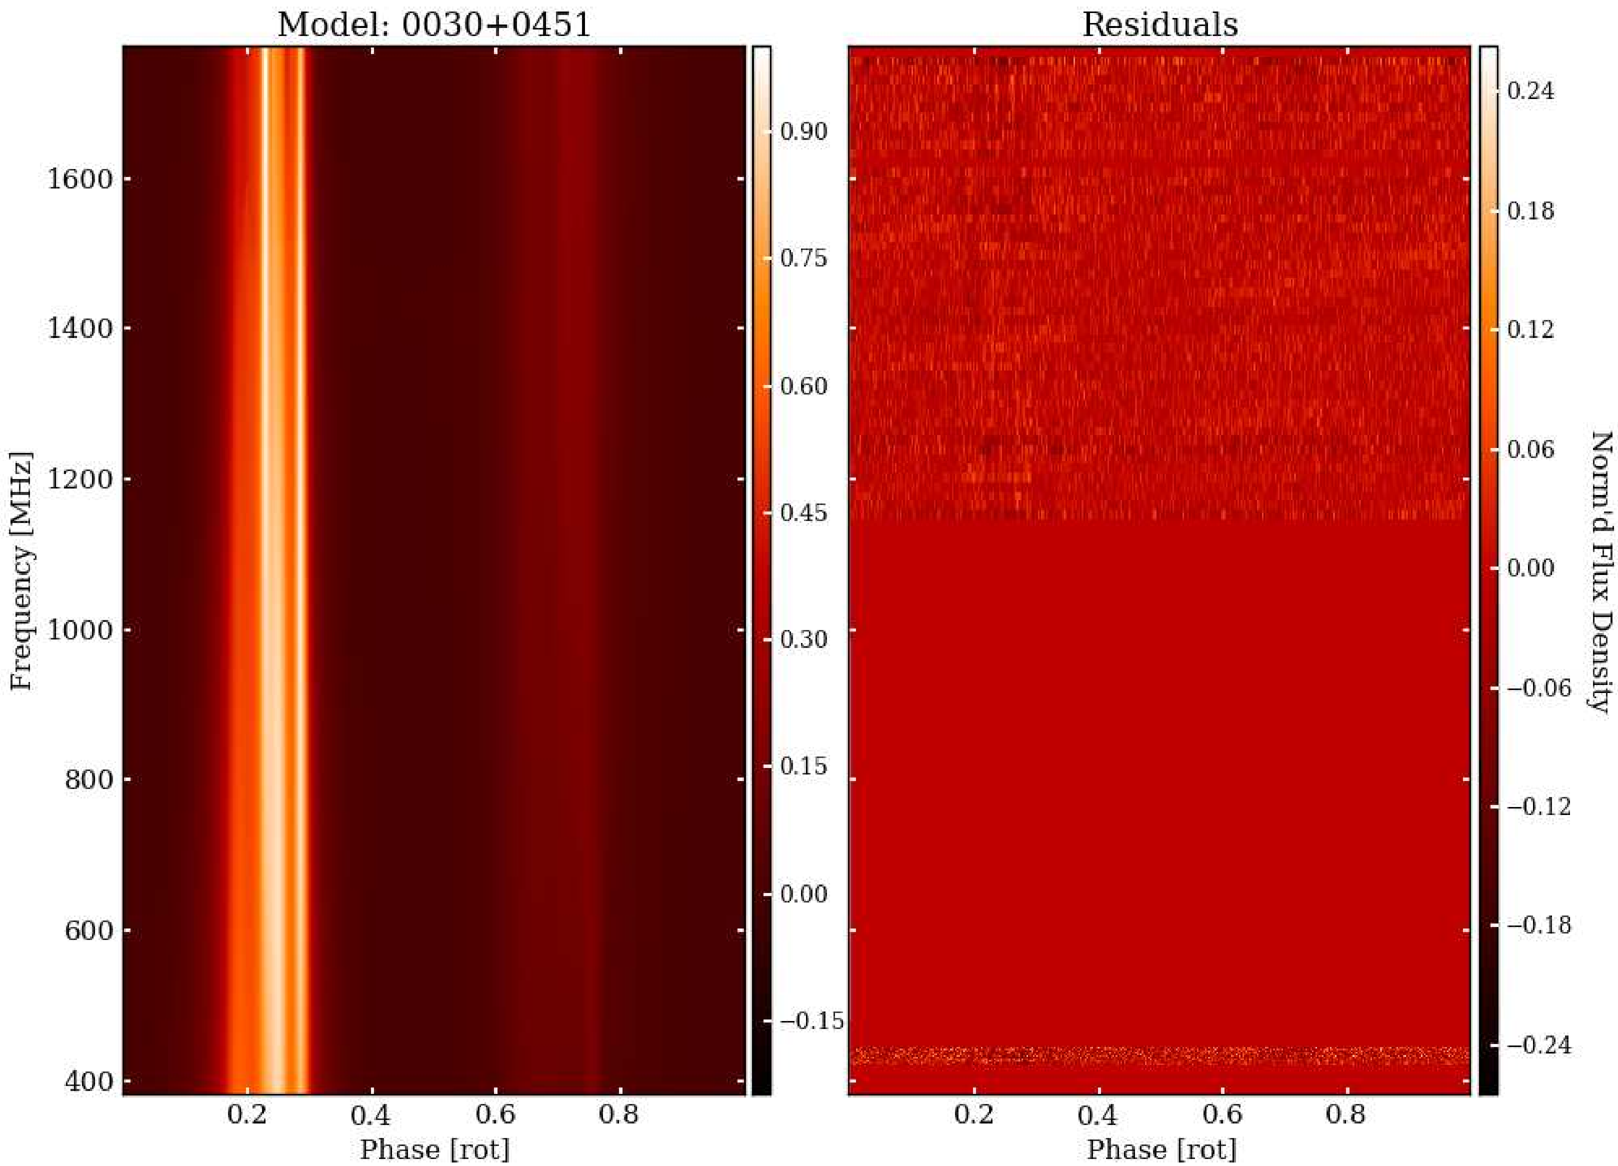

Fig. 4.159 – Fitted Gaussian component model and residuals for J0030+0451.

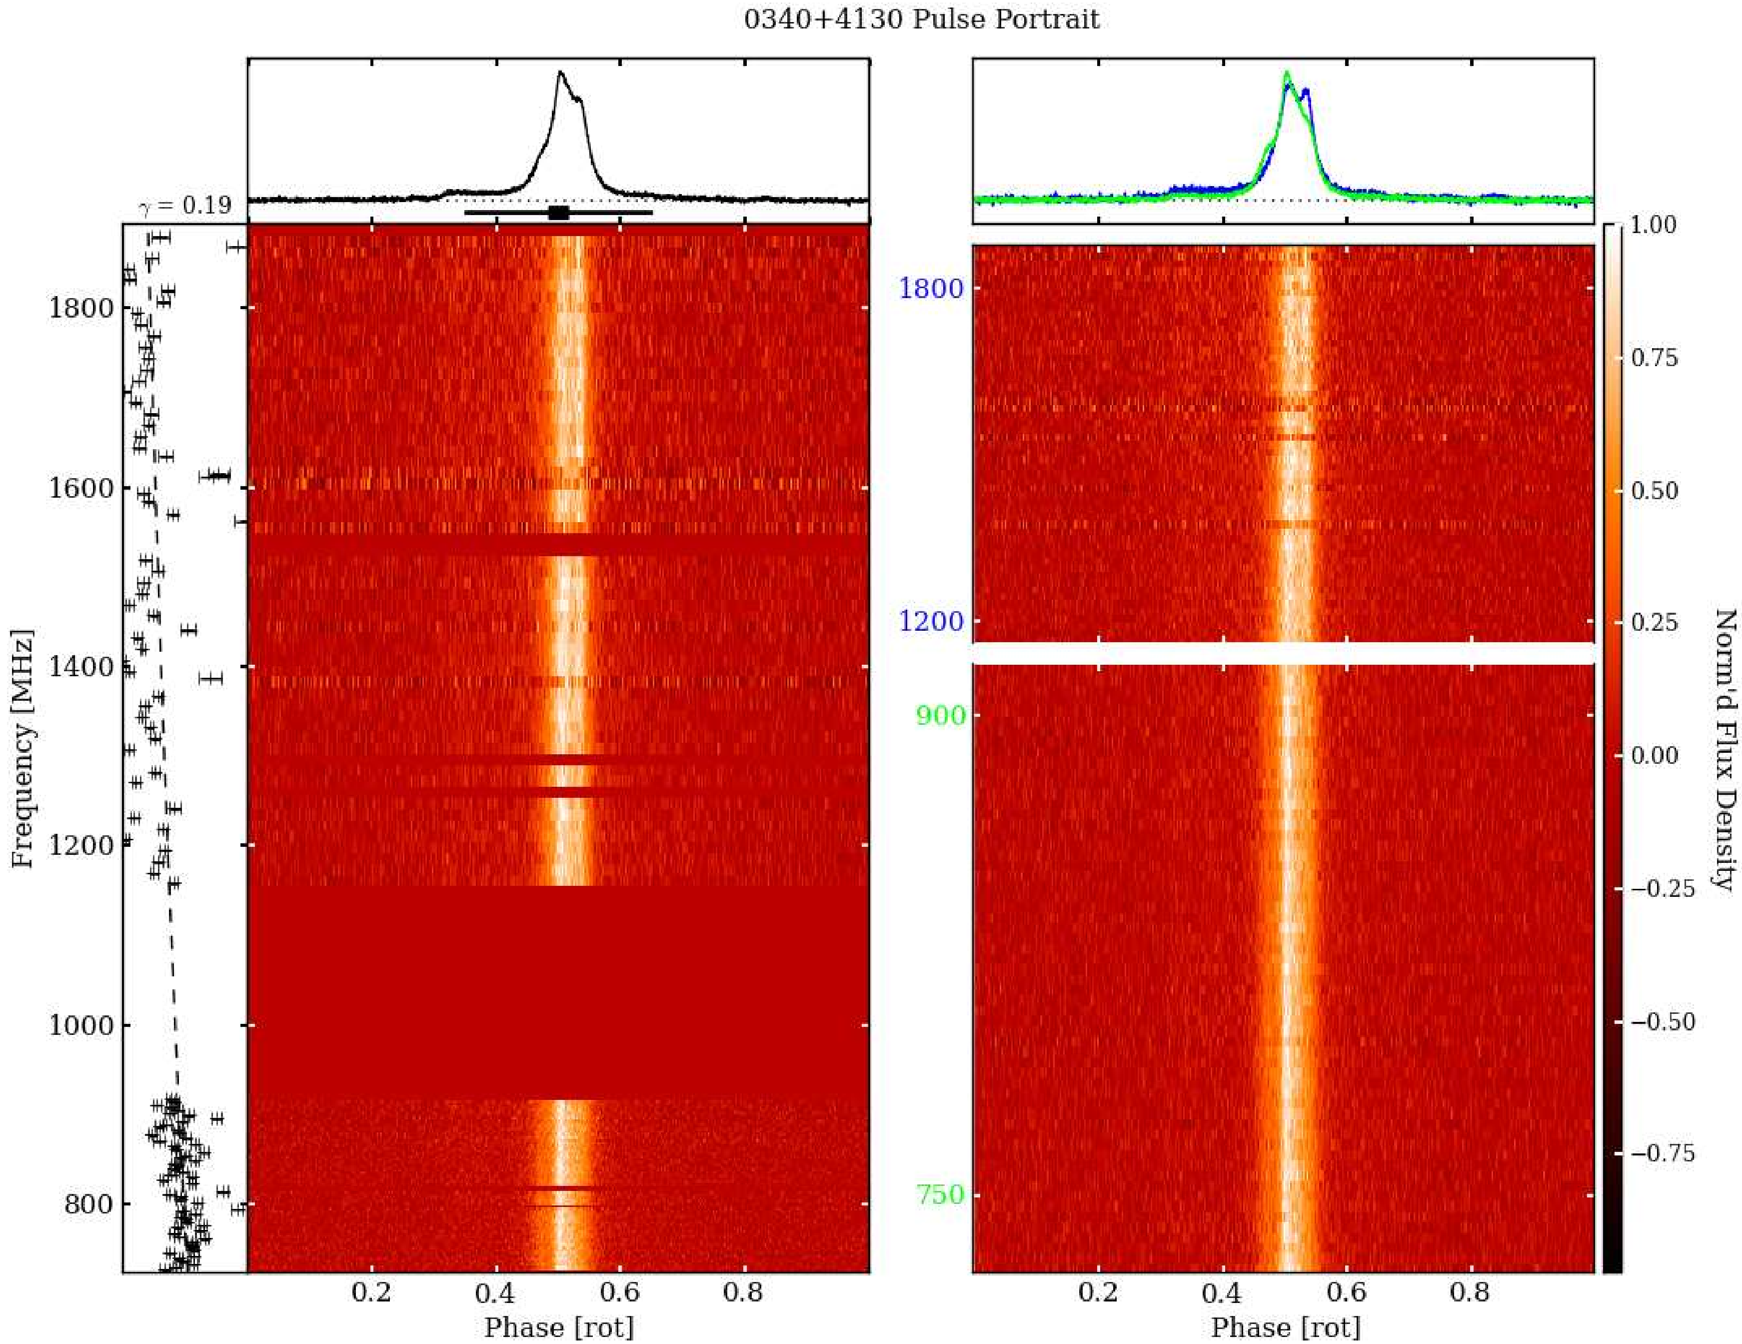

Fig. 4.160 – Concatenated portrait data and average profiles for J0340+4130.

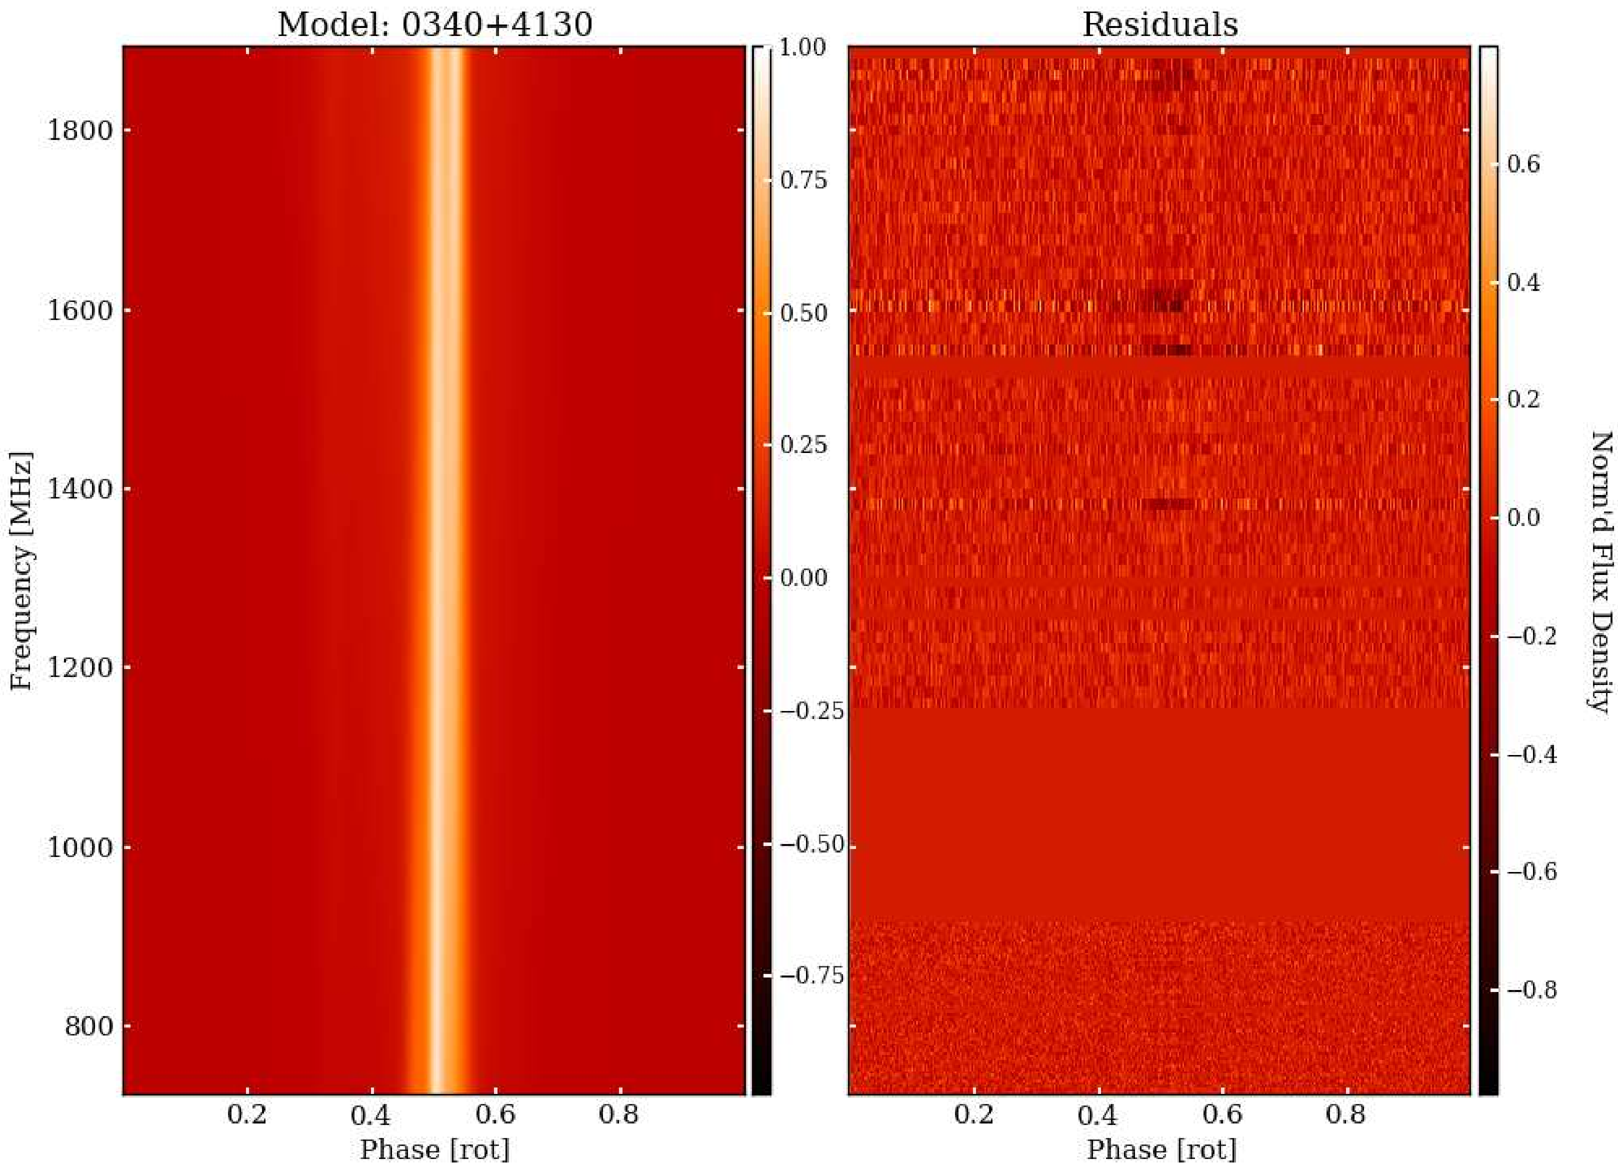

Fig. 4.161 – Fitted Gaussian component model and residuals for J0340+4130.

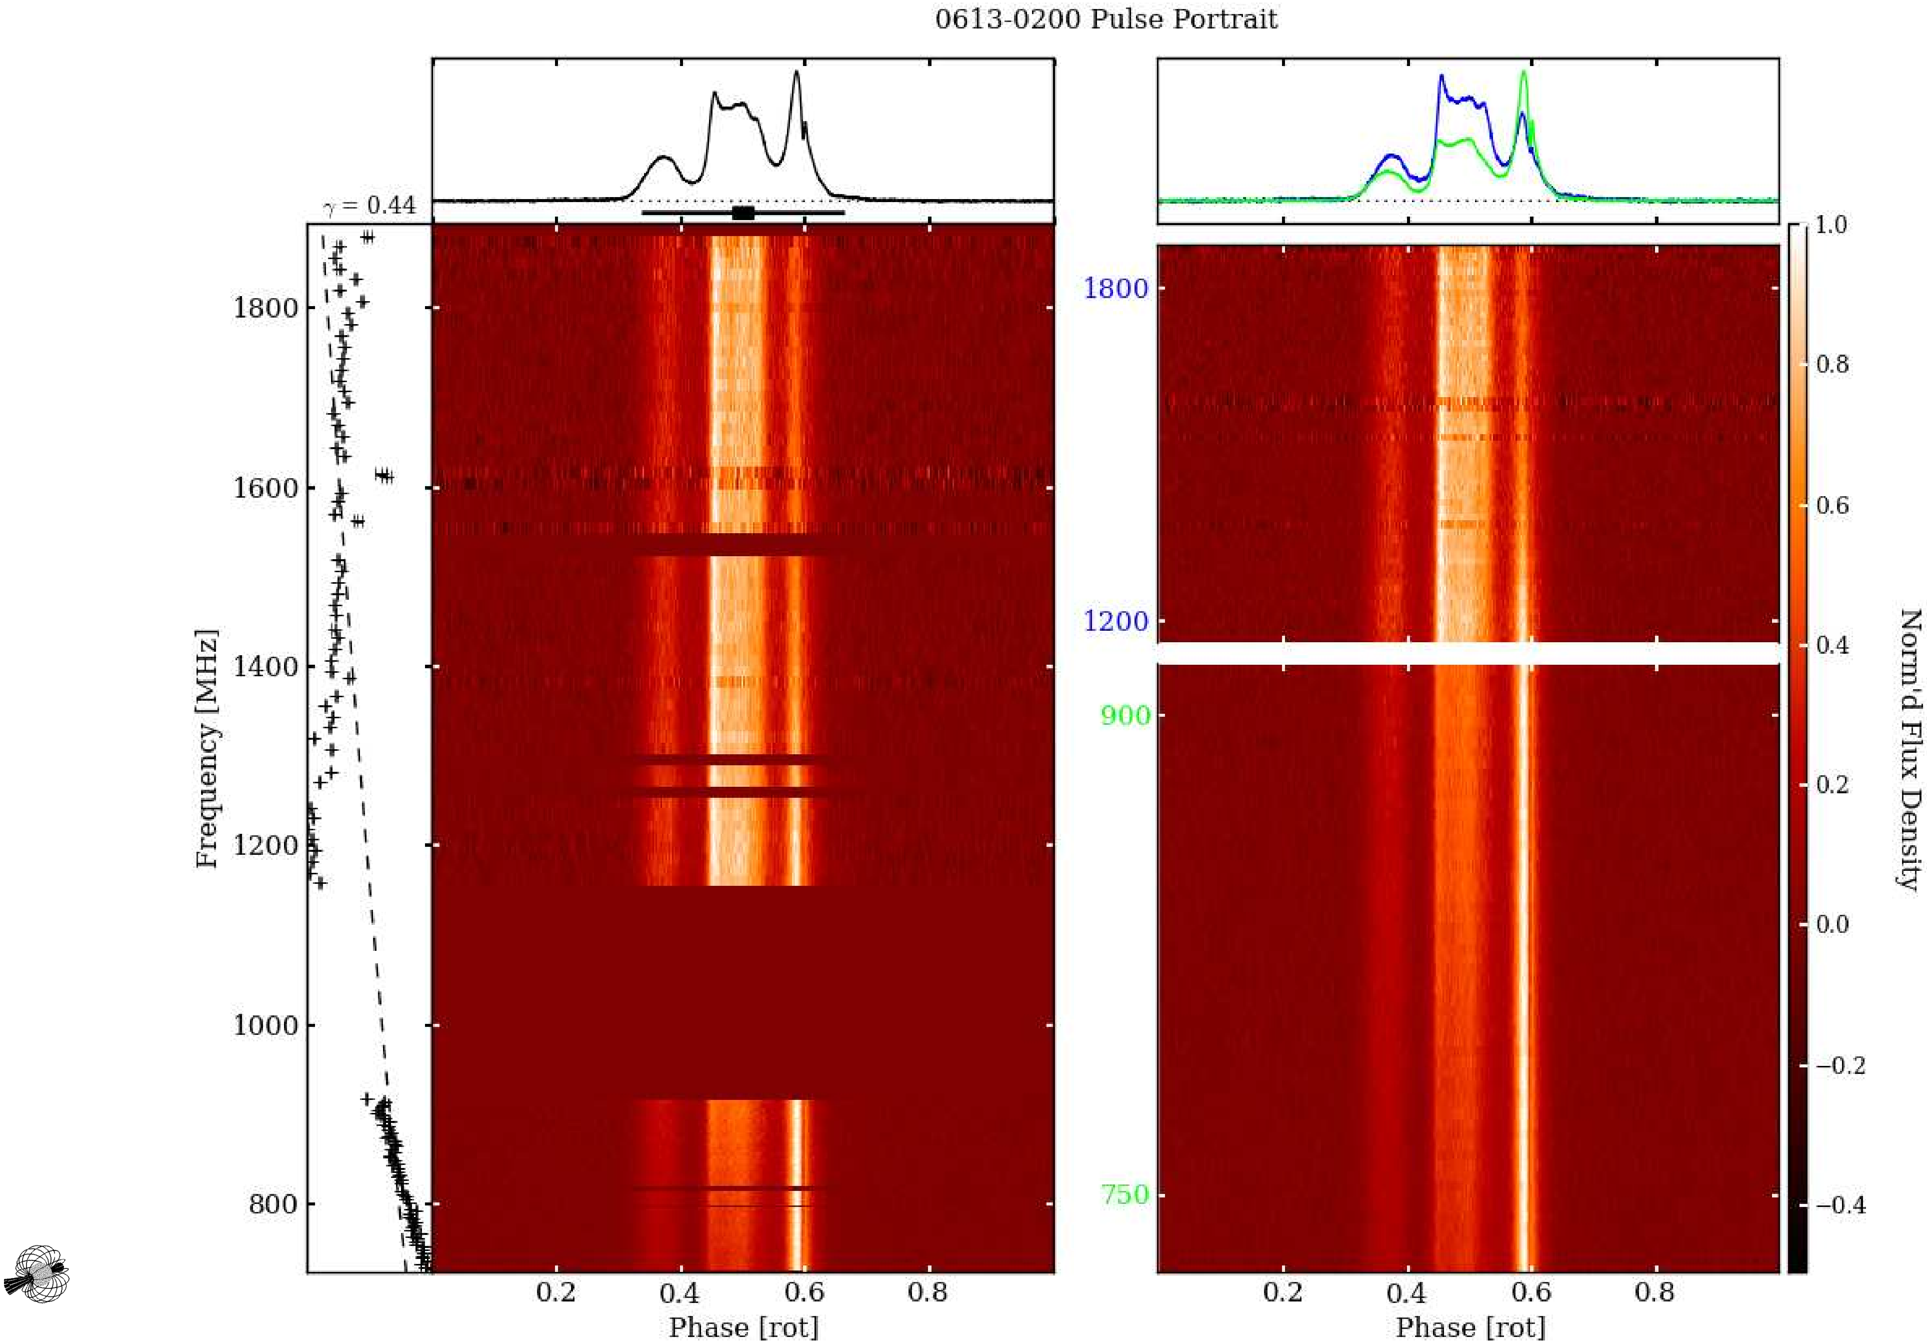

Fig. 4.162 – Concatenated portrait data and average profiles for J0613−0200.

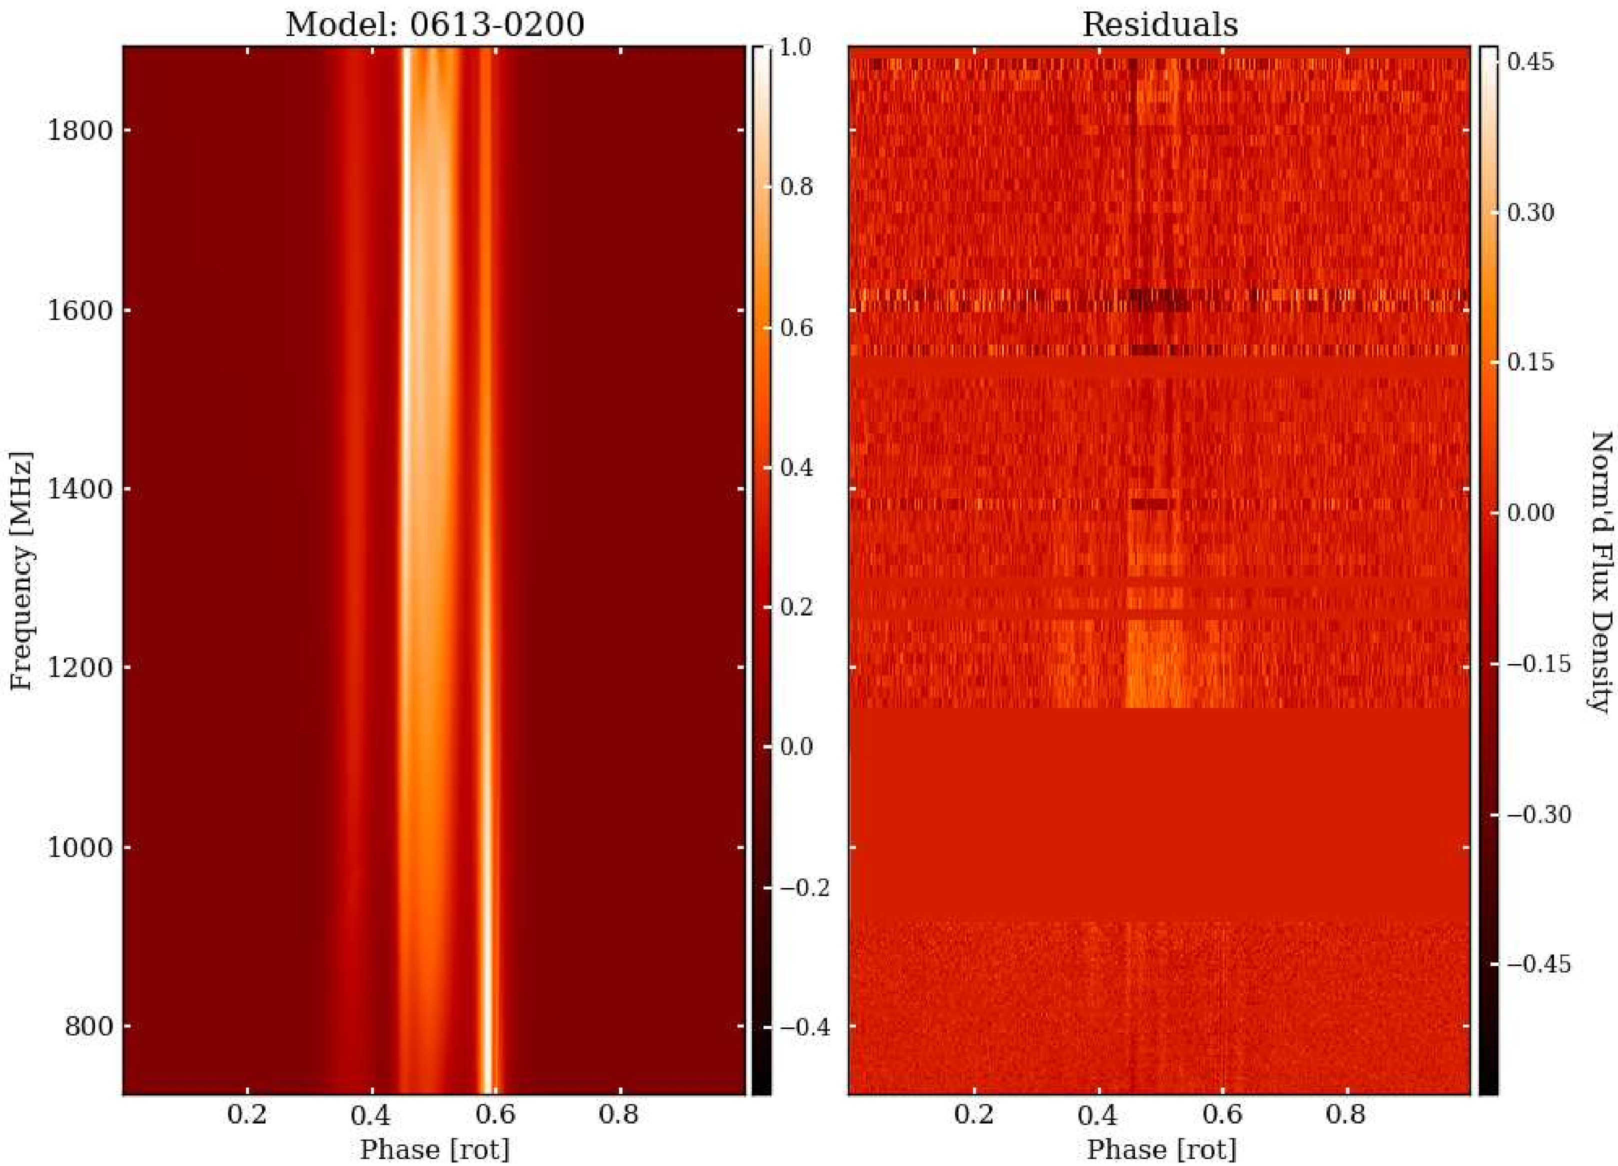

Fig. 4.163 – Fitted Gaussian component model and residuals for J0613−0200.

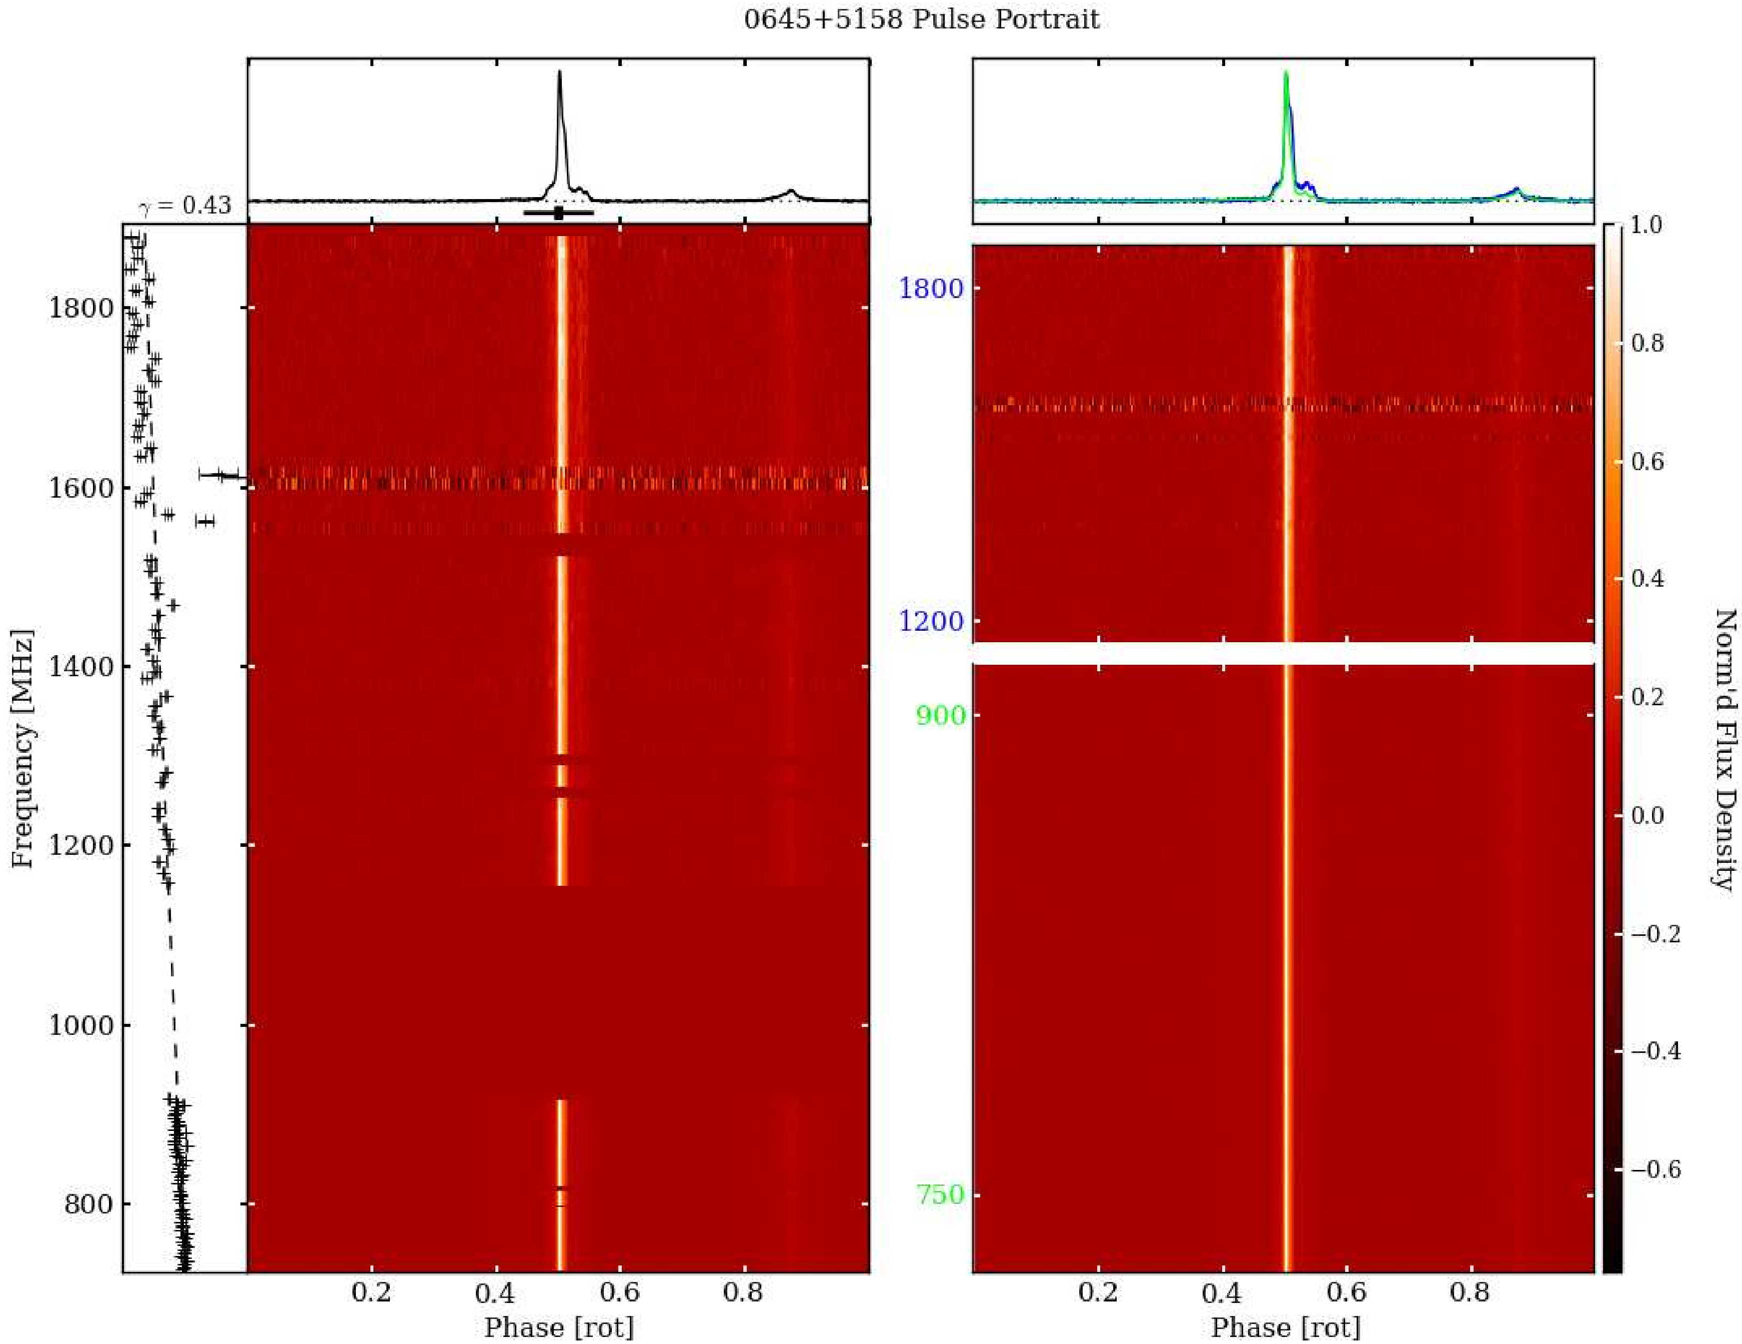

Fig. 4.164 – Concatenated portrait data and average profiles for J0645+5158.

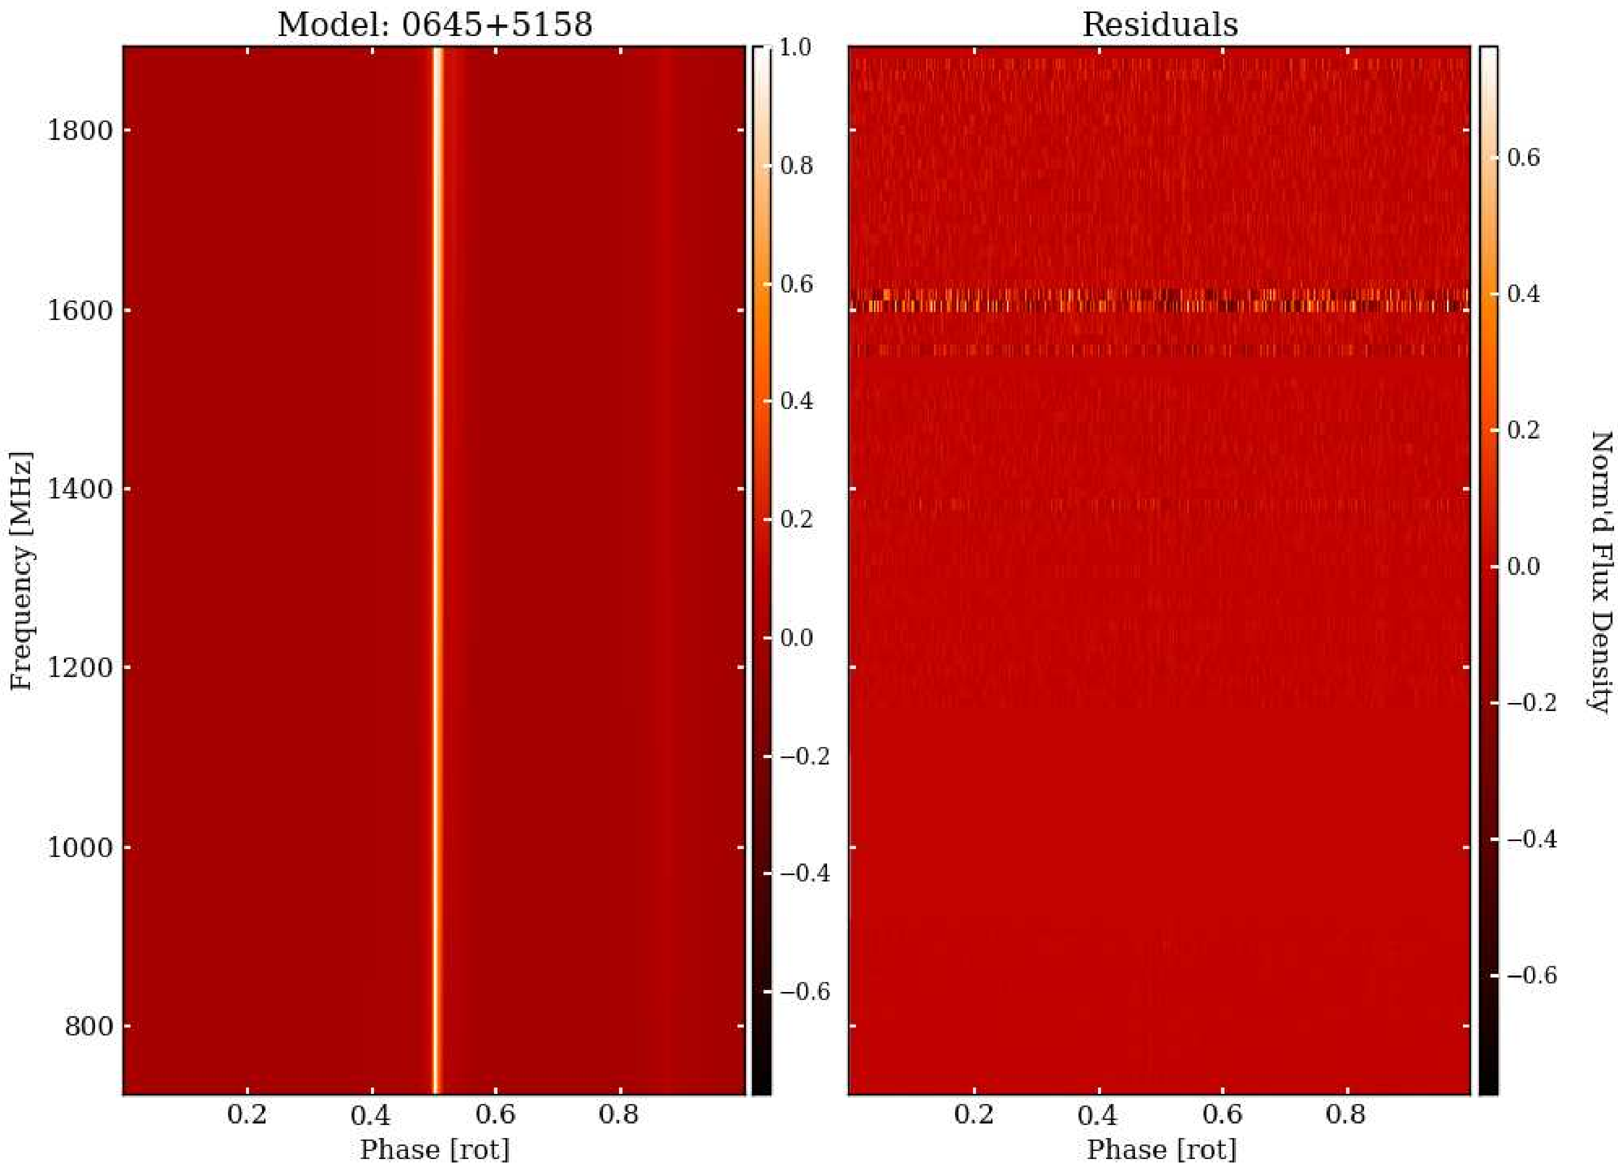

Fig. 4.165 – Fitted Gaussian component model and residuals for J0645+5158.

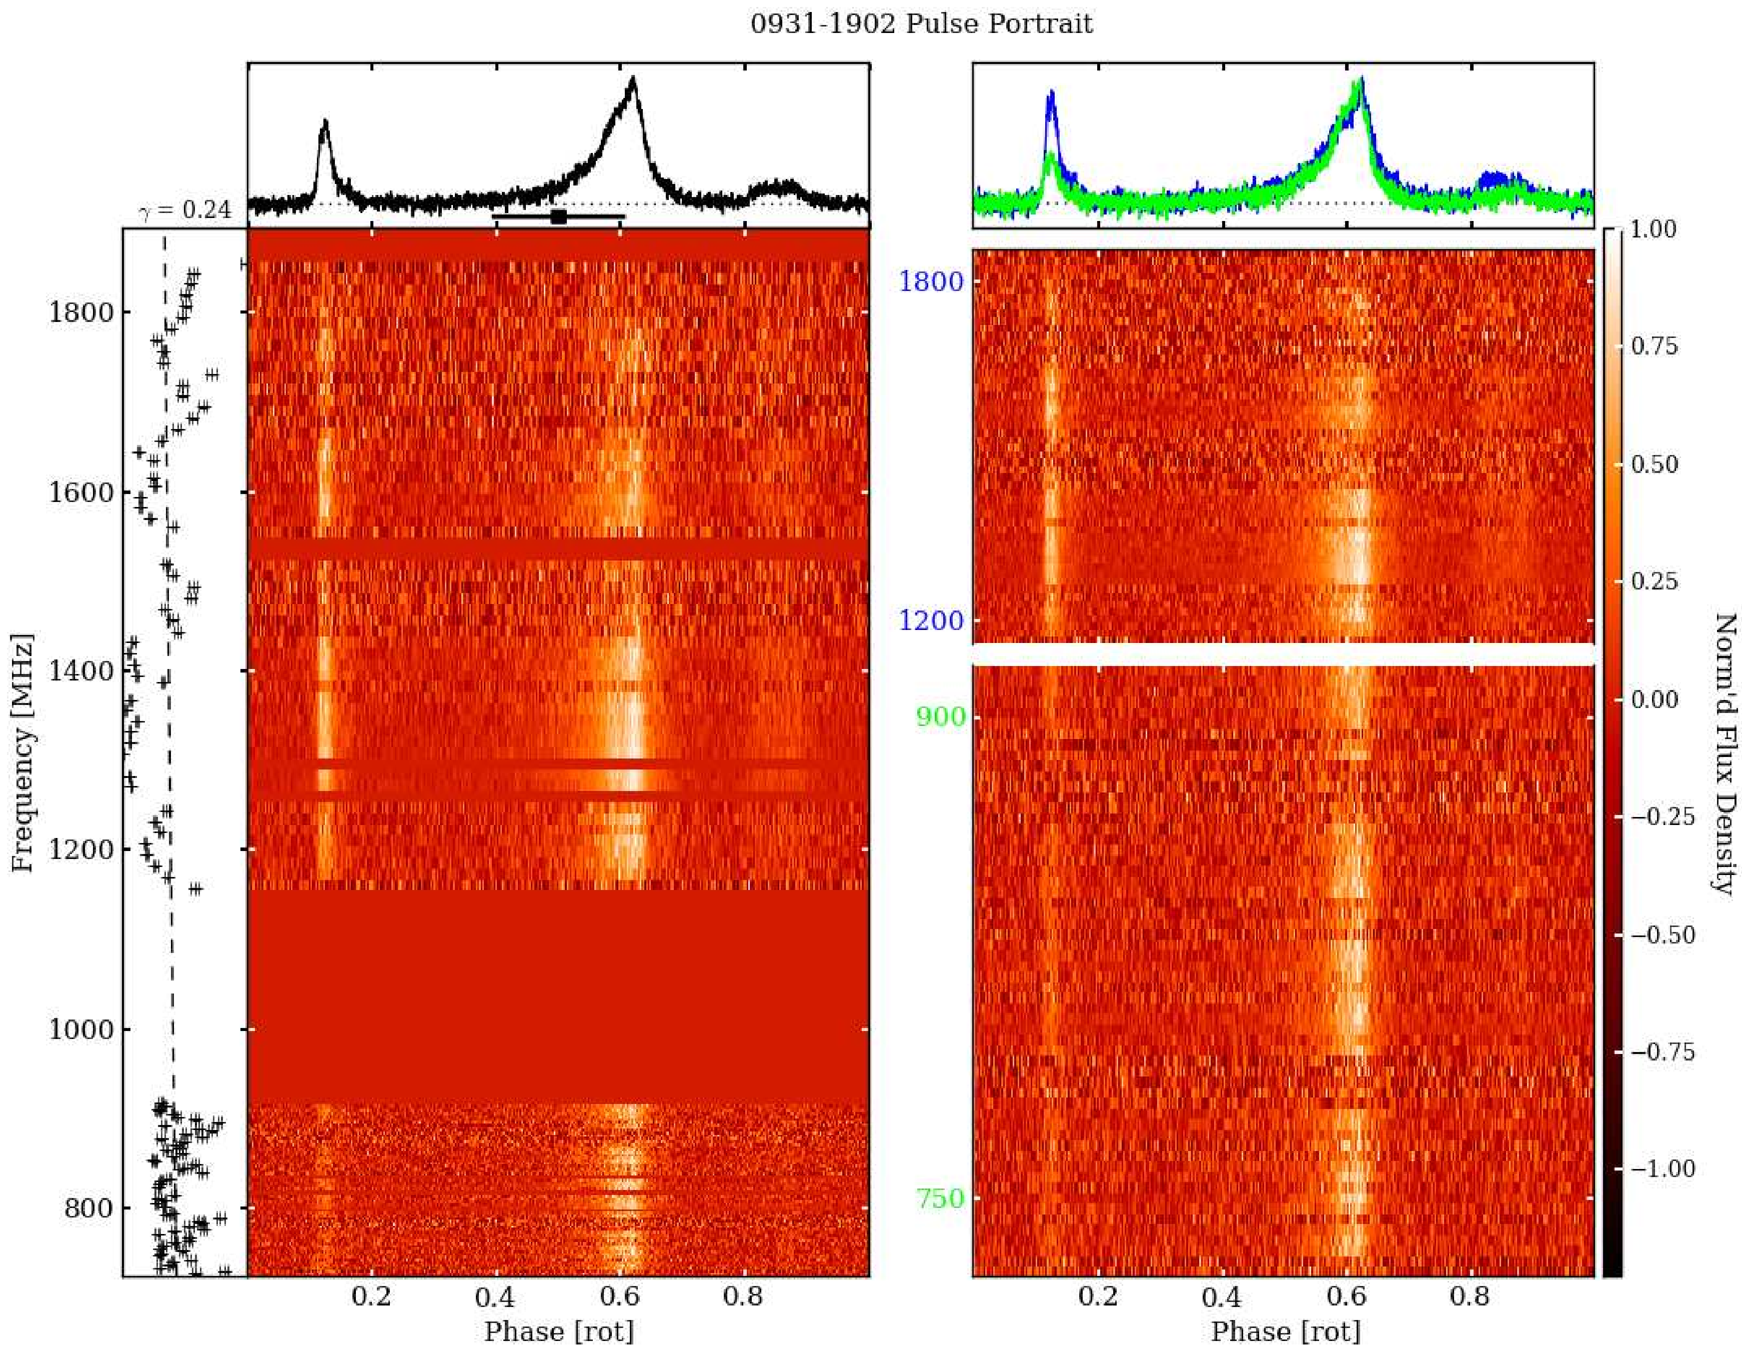

Fig. 4.166 – Concatenated portrait data and average profiles for J0931−1902.

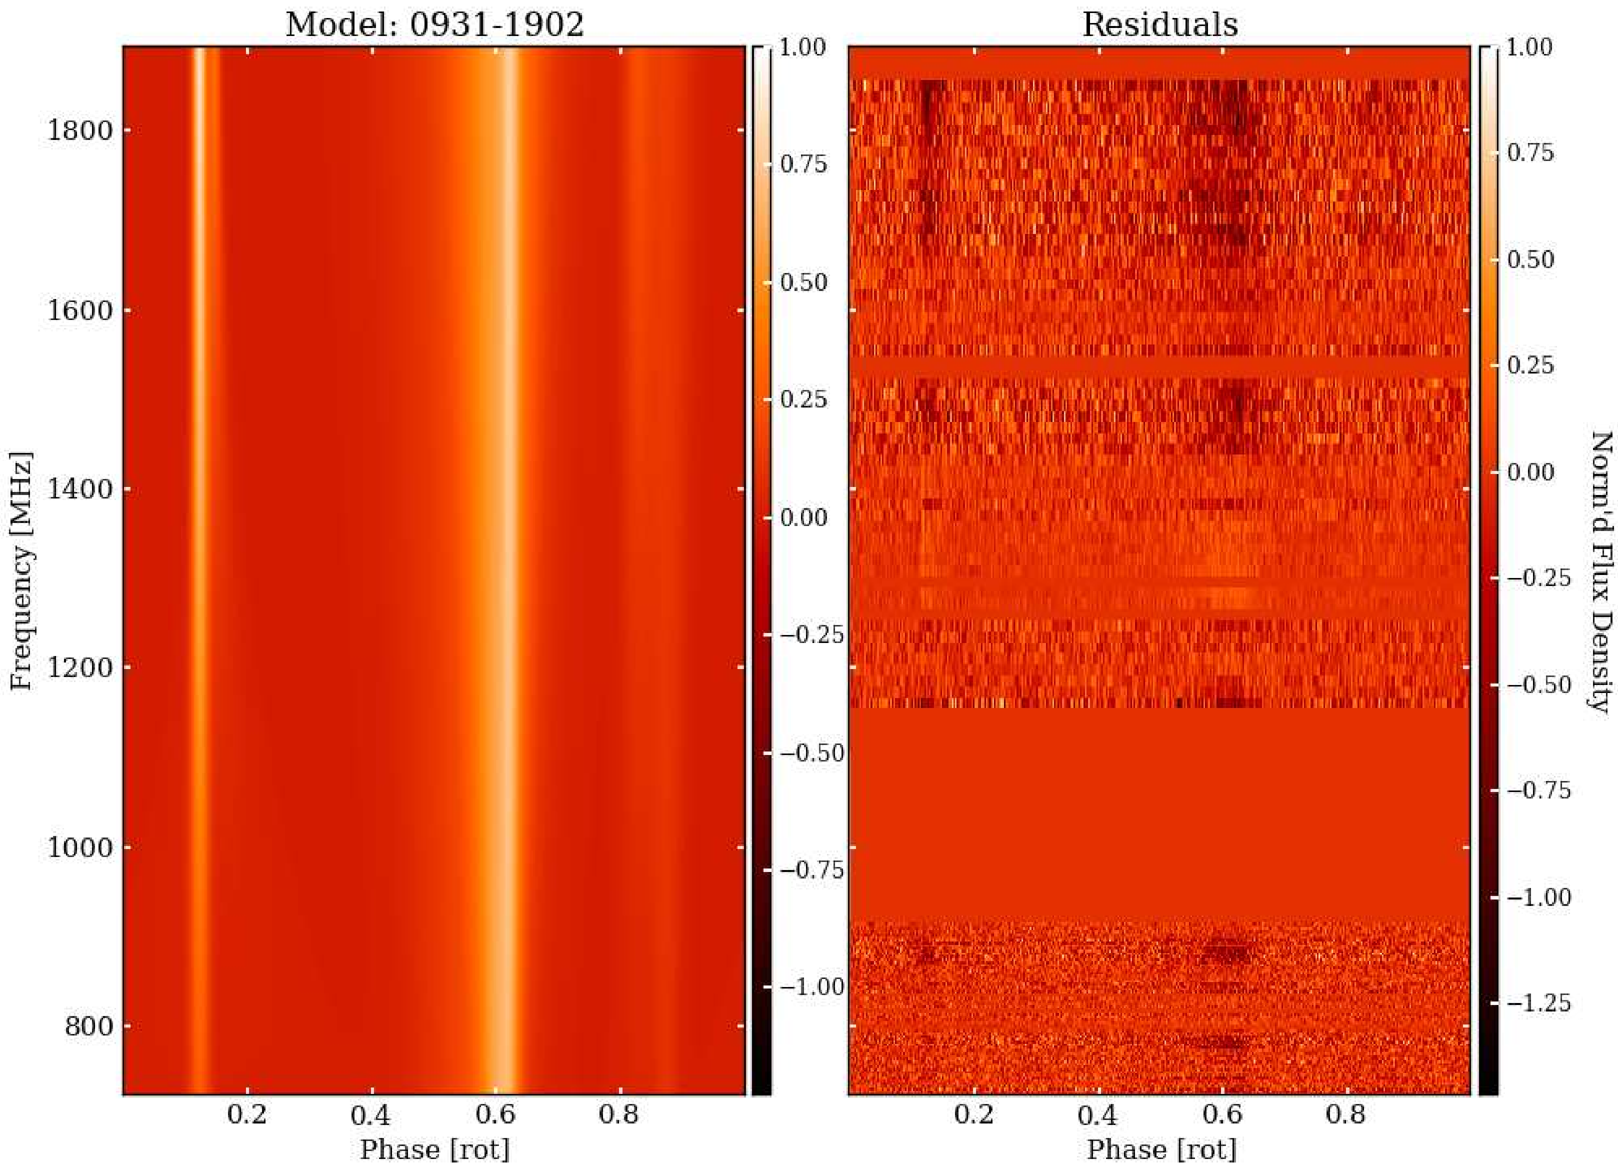

Fig. 4.167 – Fitted Gaussian component model and residuals for J0931−1902.

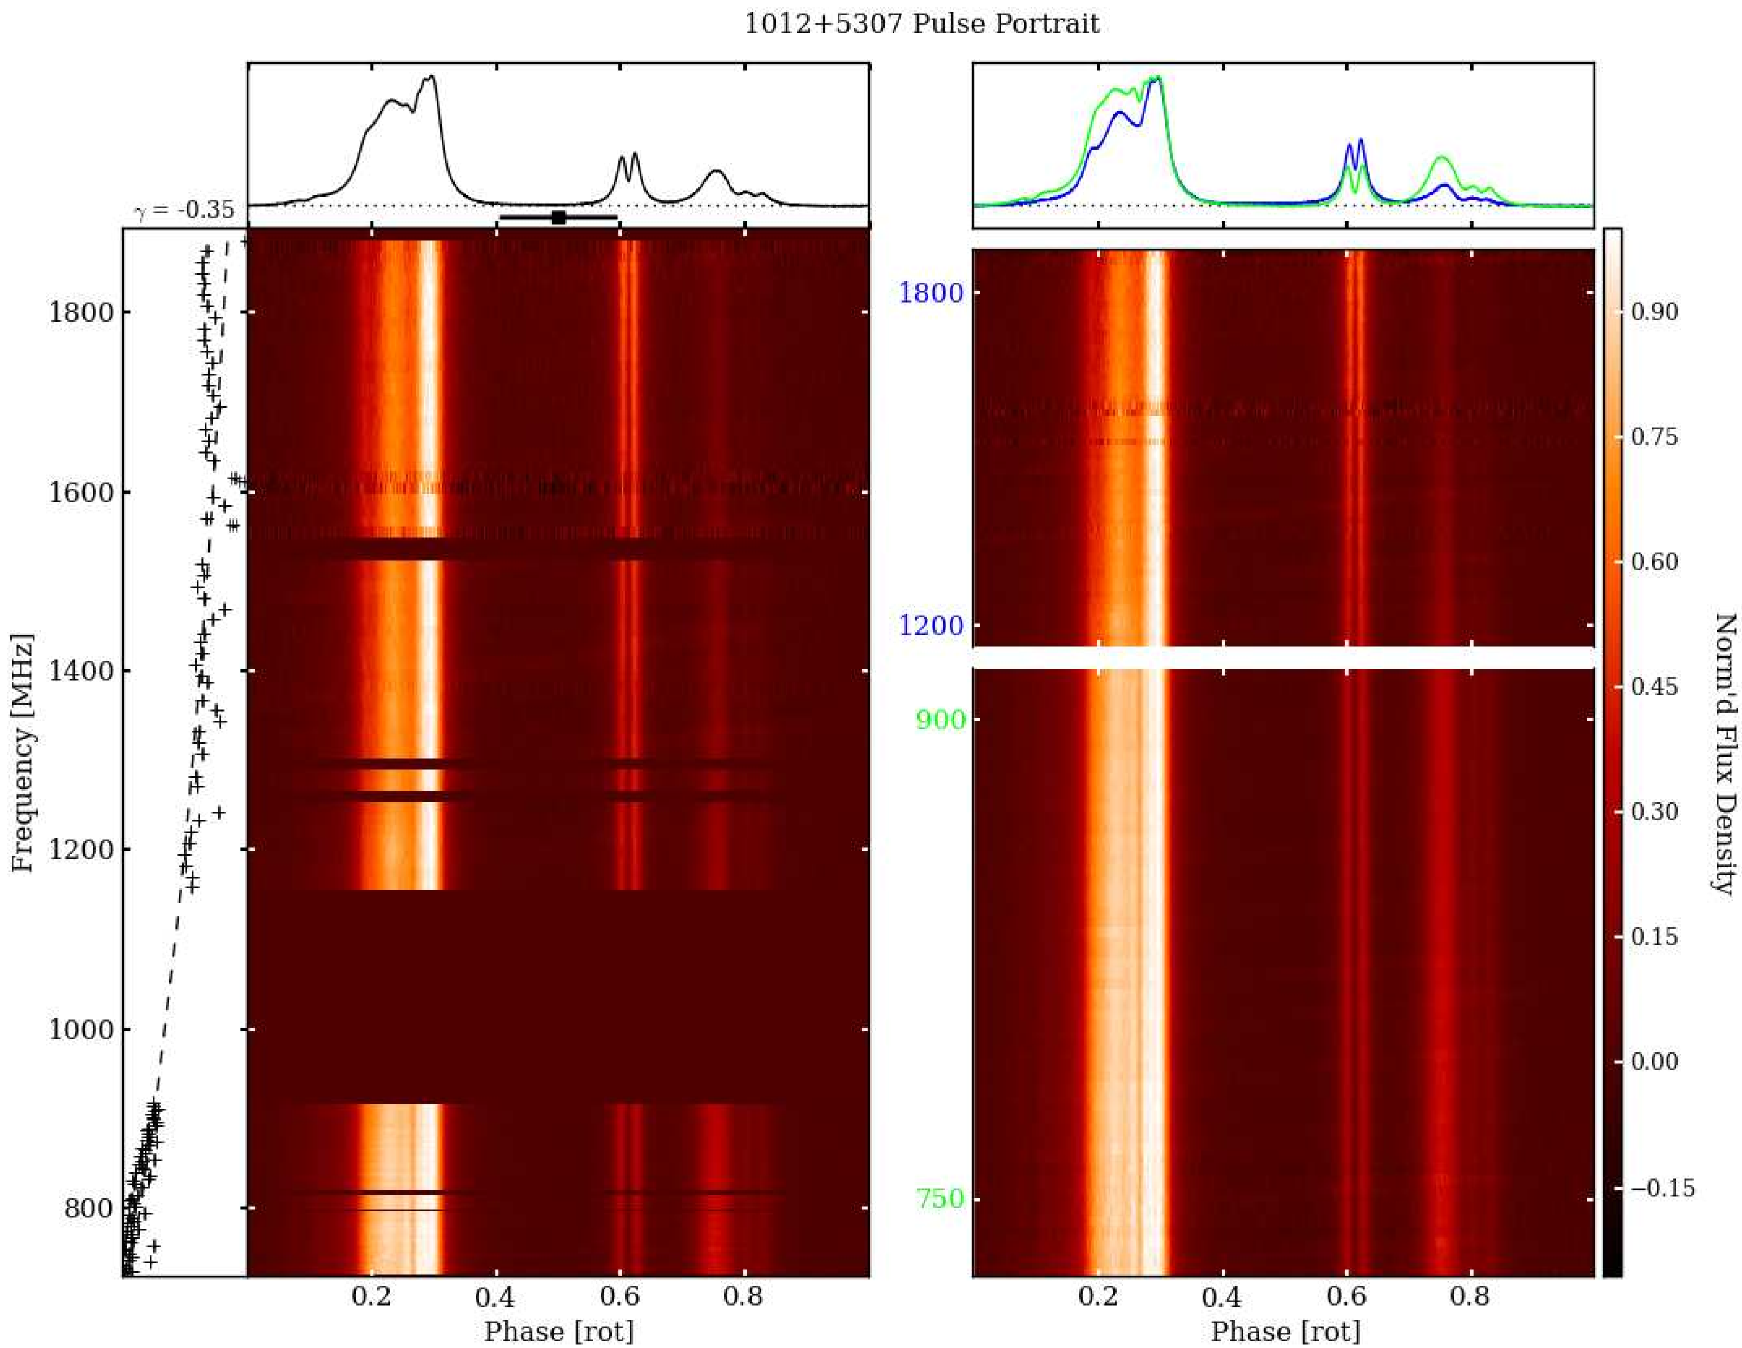

Fig. 4.168 – Concatenated portrait data and average profiles for J1012+5307.

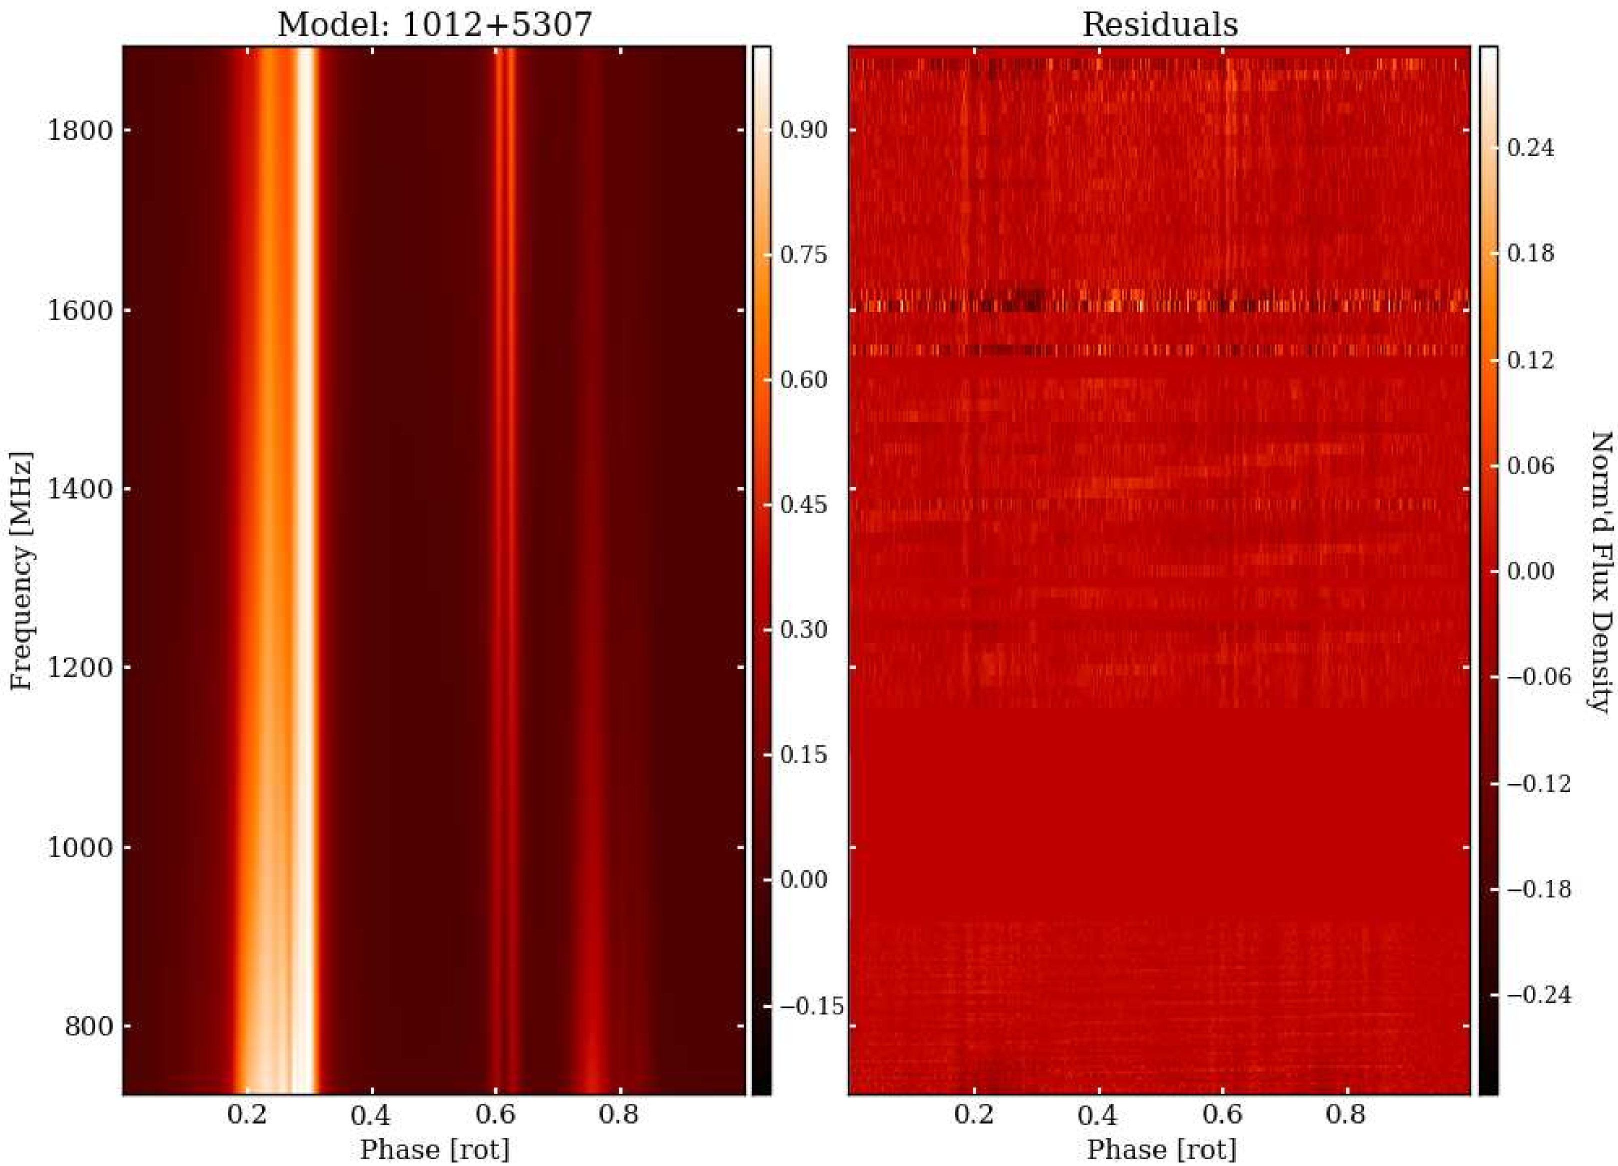

Fig. 4.169 – Fitted Gaussian component model and residuals for J1012+5307.

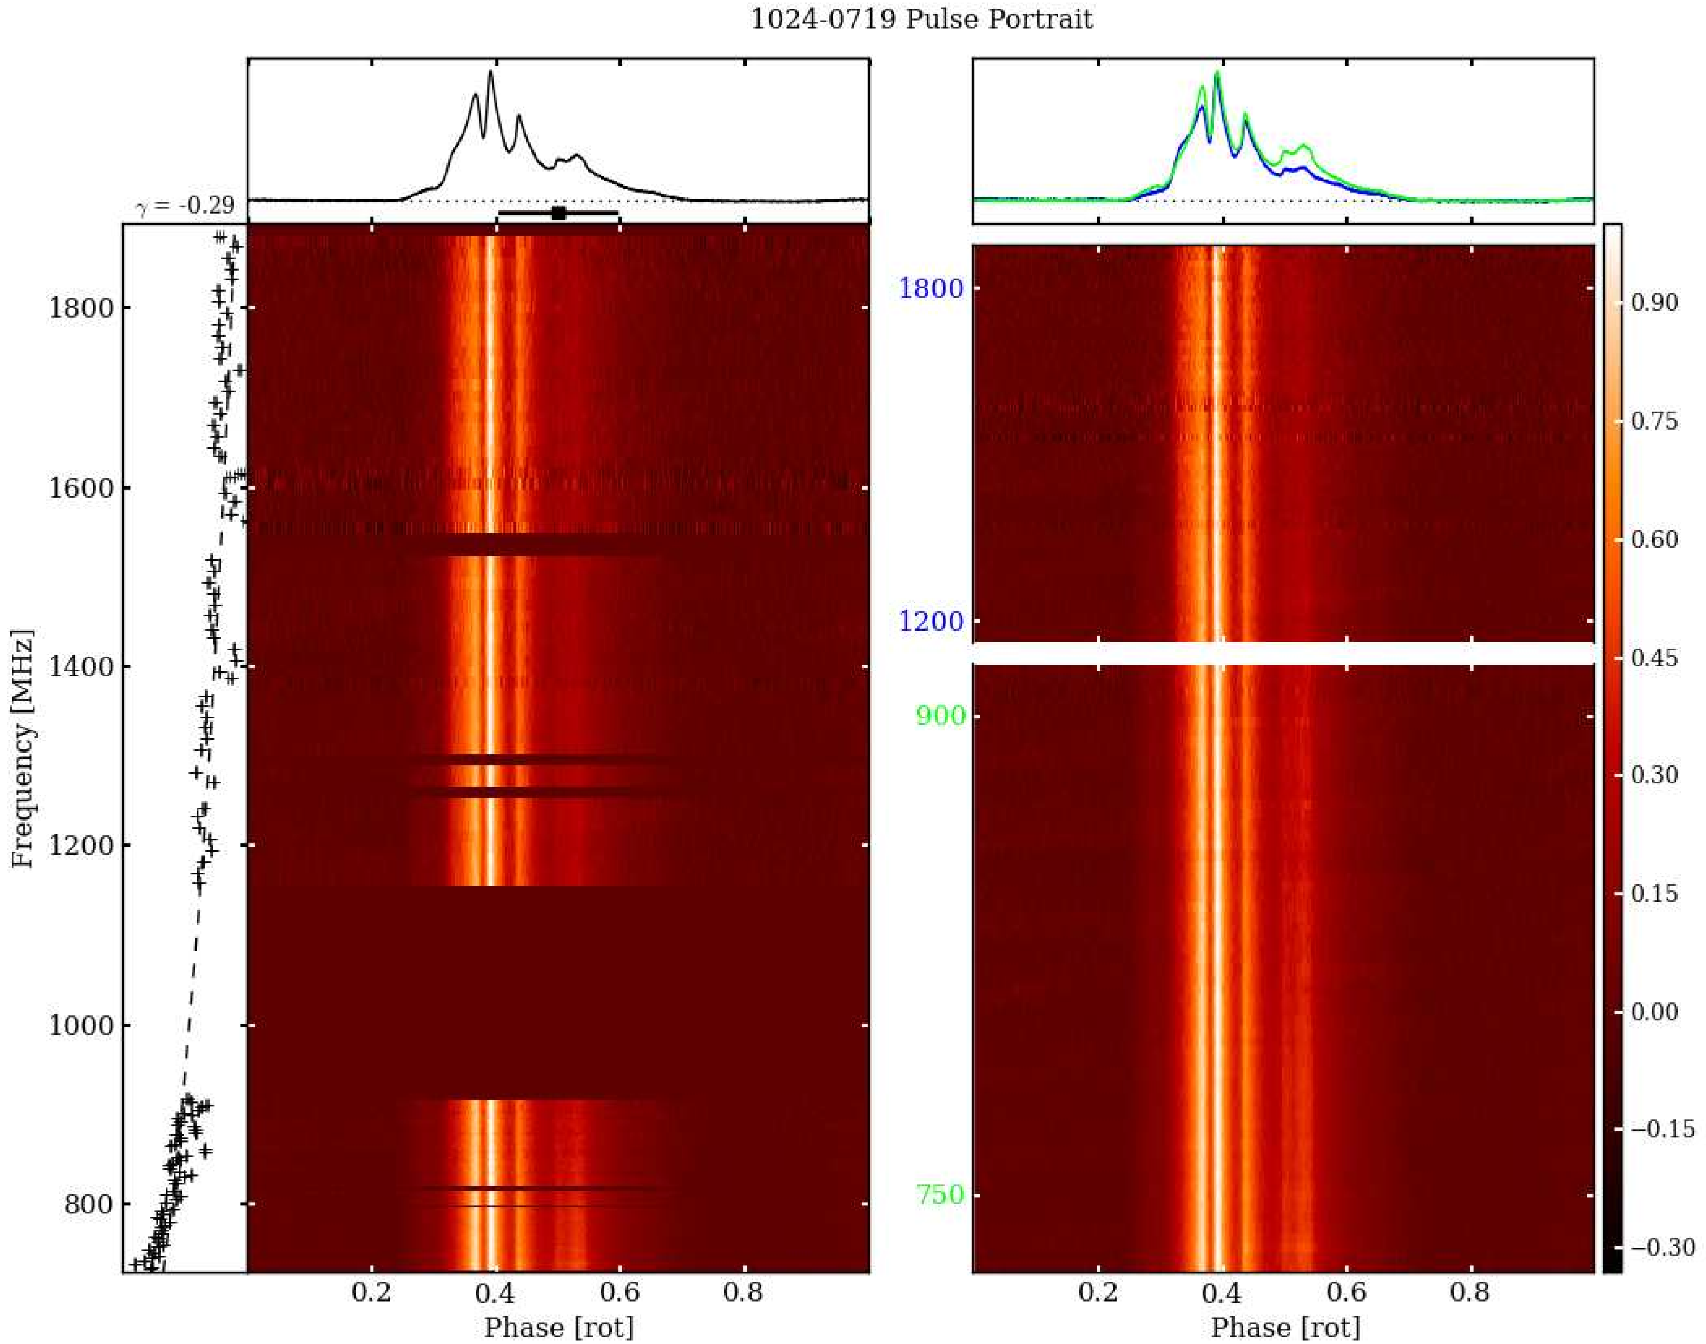

Fig. 4.170 – Concatenated portrait data and average profiles for J1024−0719.

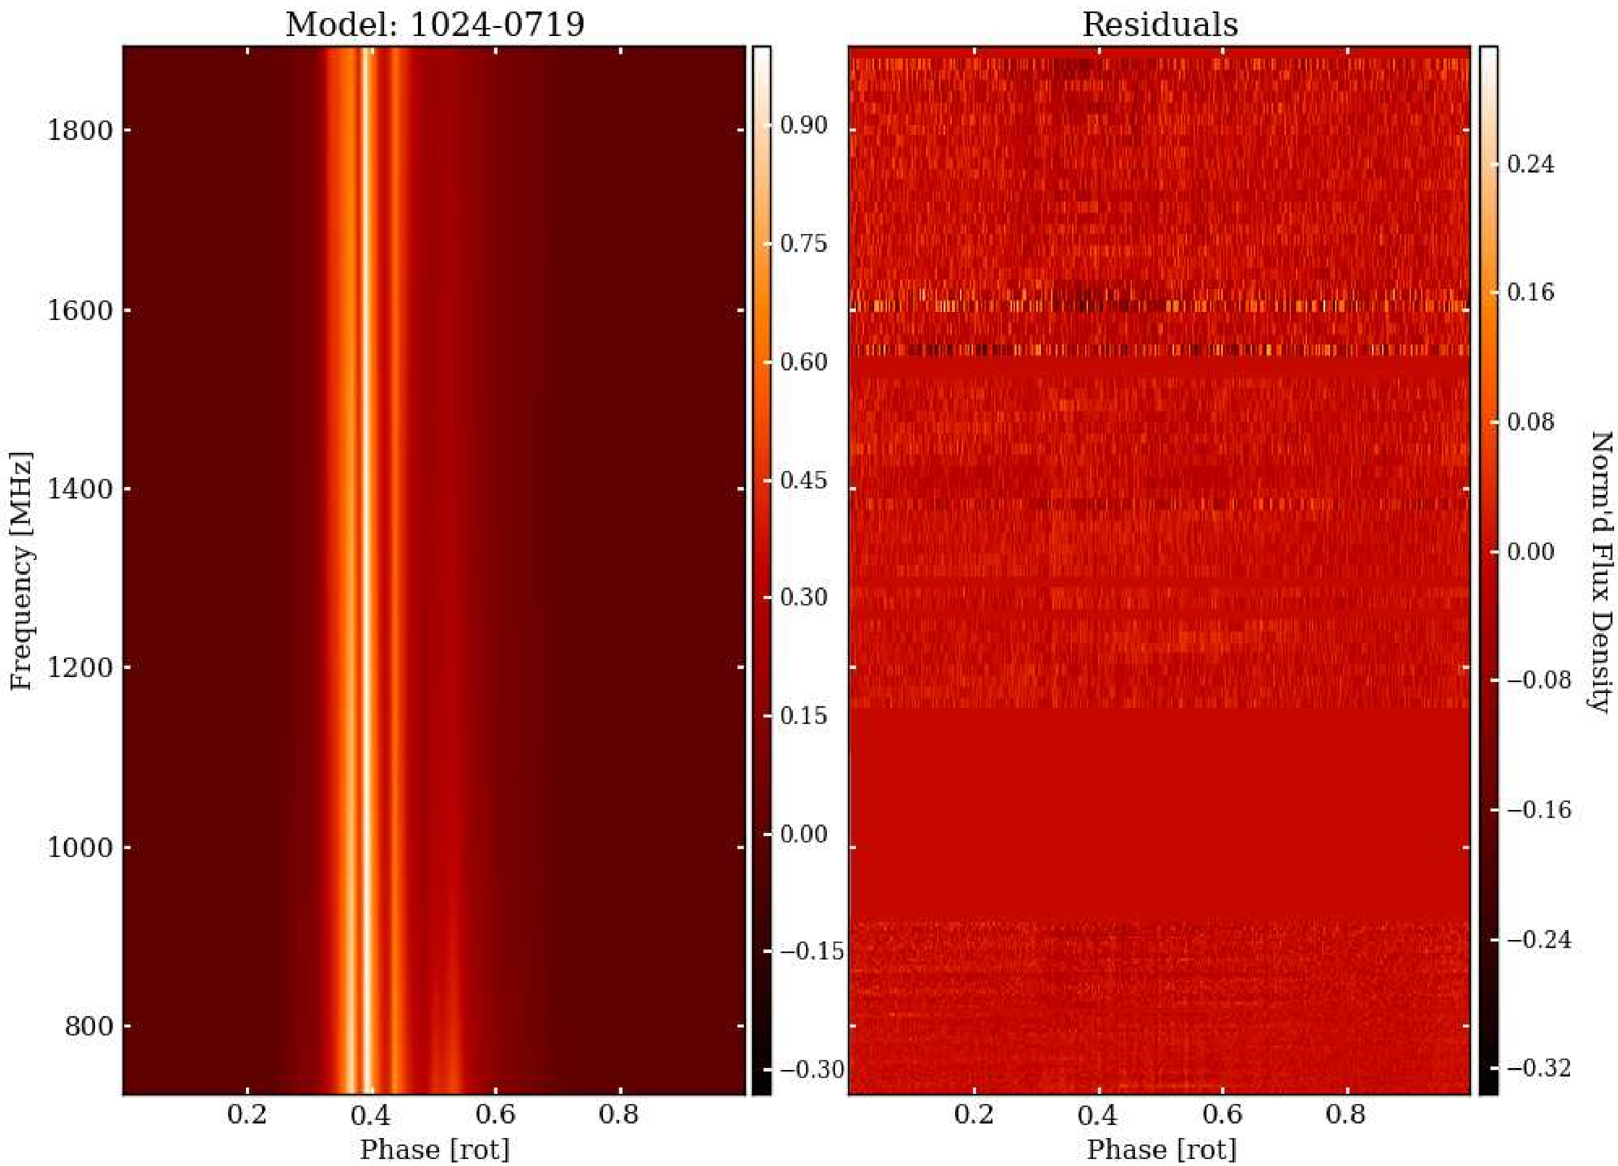

Fig. 4.171 – Fitted Gaussian component model and residuals for J1024−0719.

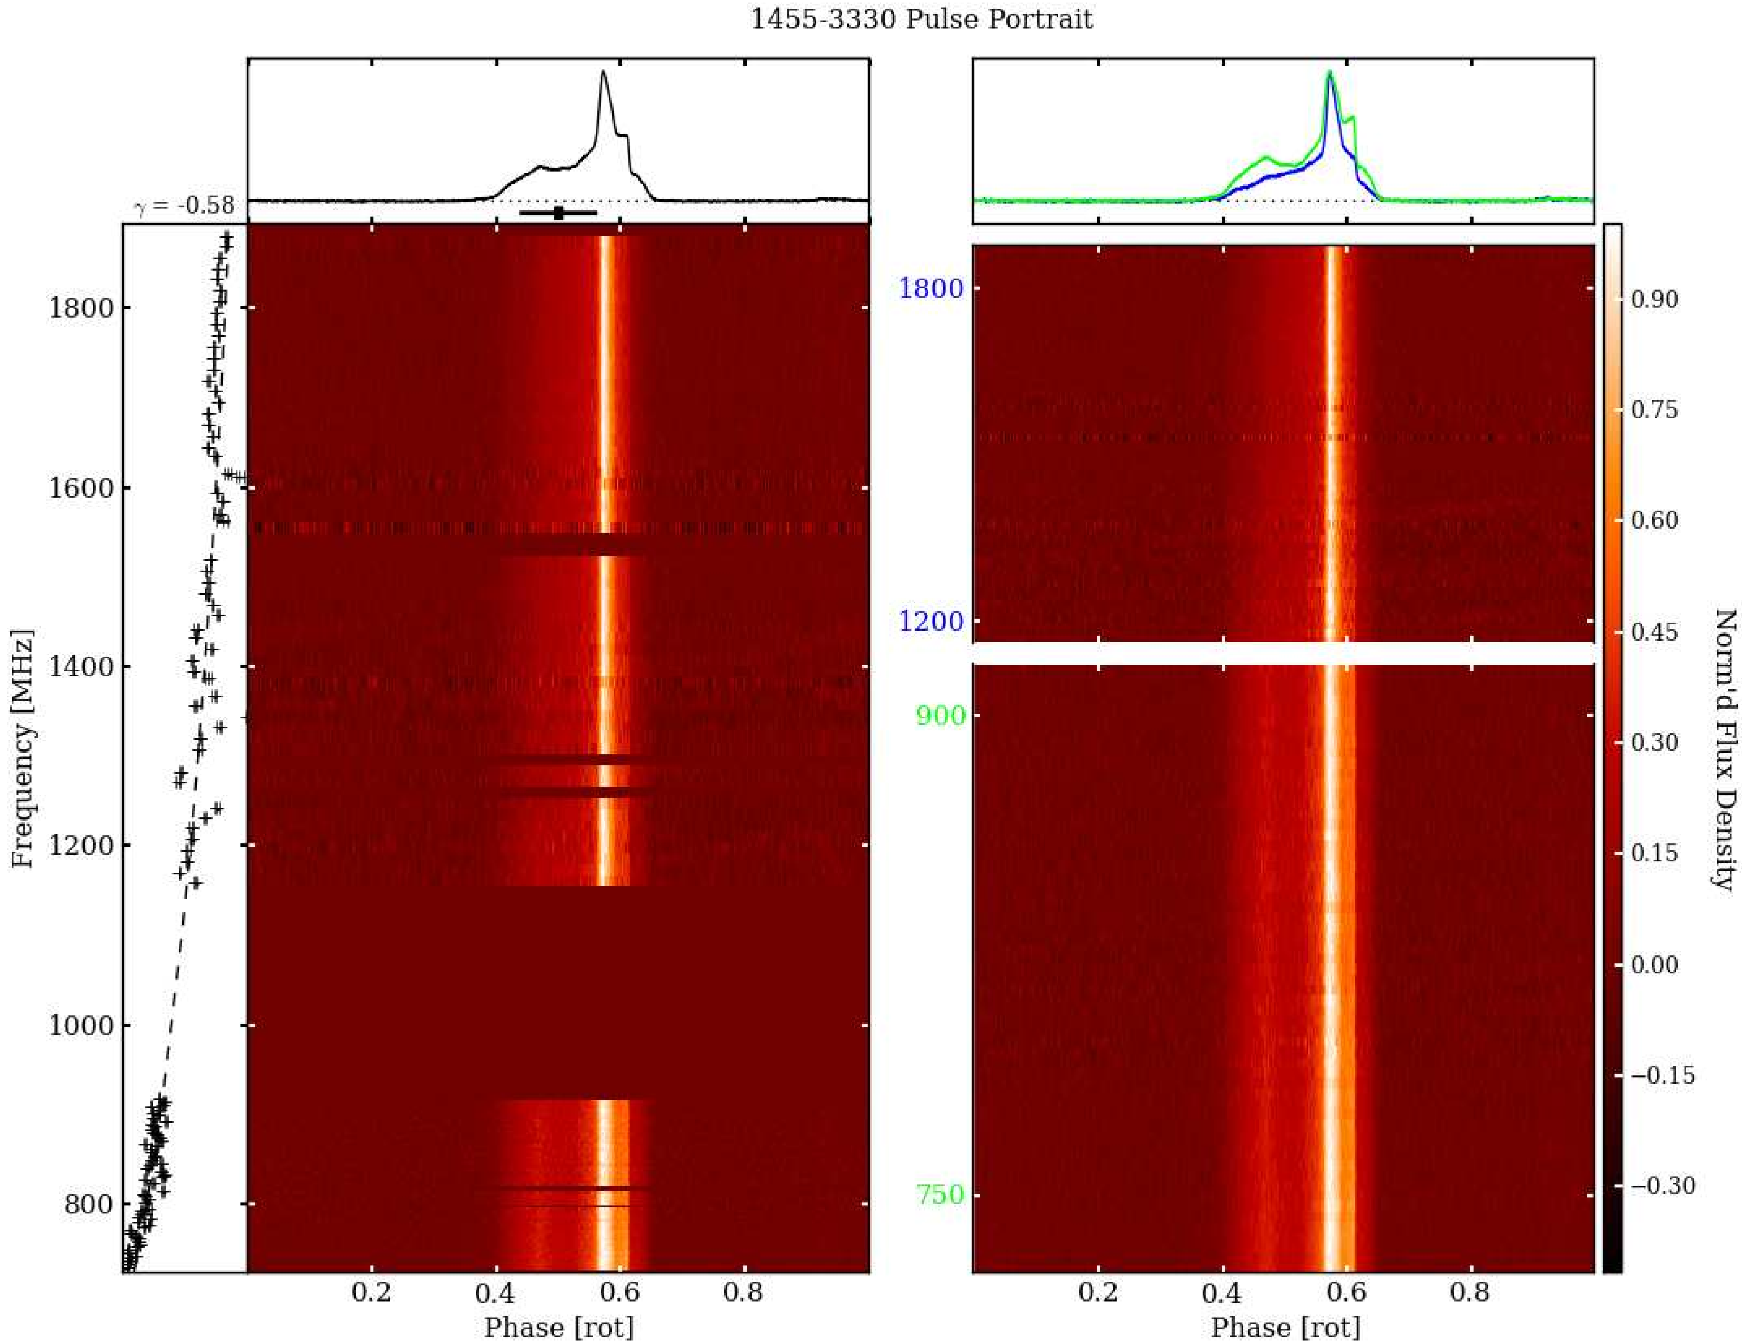

Fig. 4.172 – Concatenated portrait data and average profiles for J1455−3330.

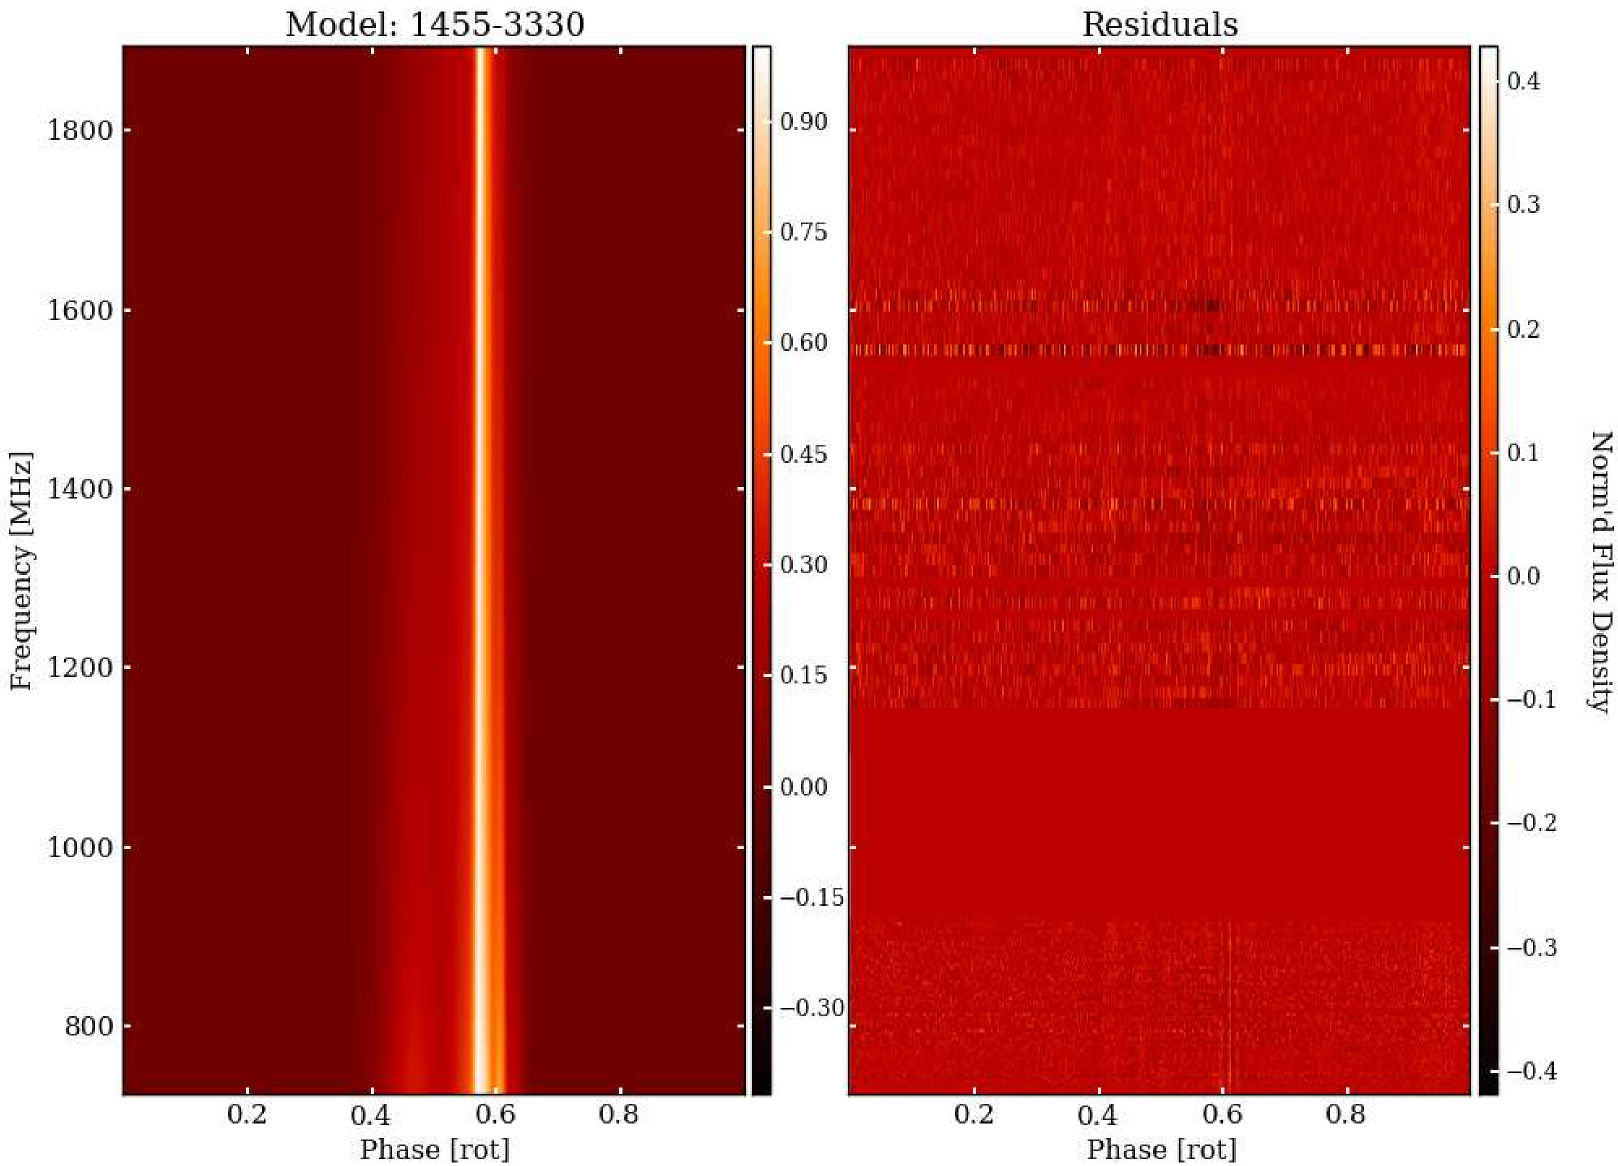

Fig. 4.173 – Fitted Gaussian component model and residuals for J1455−3330.

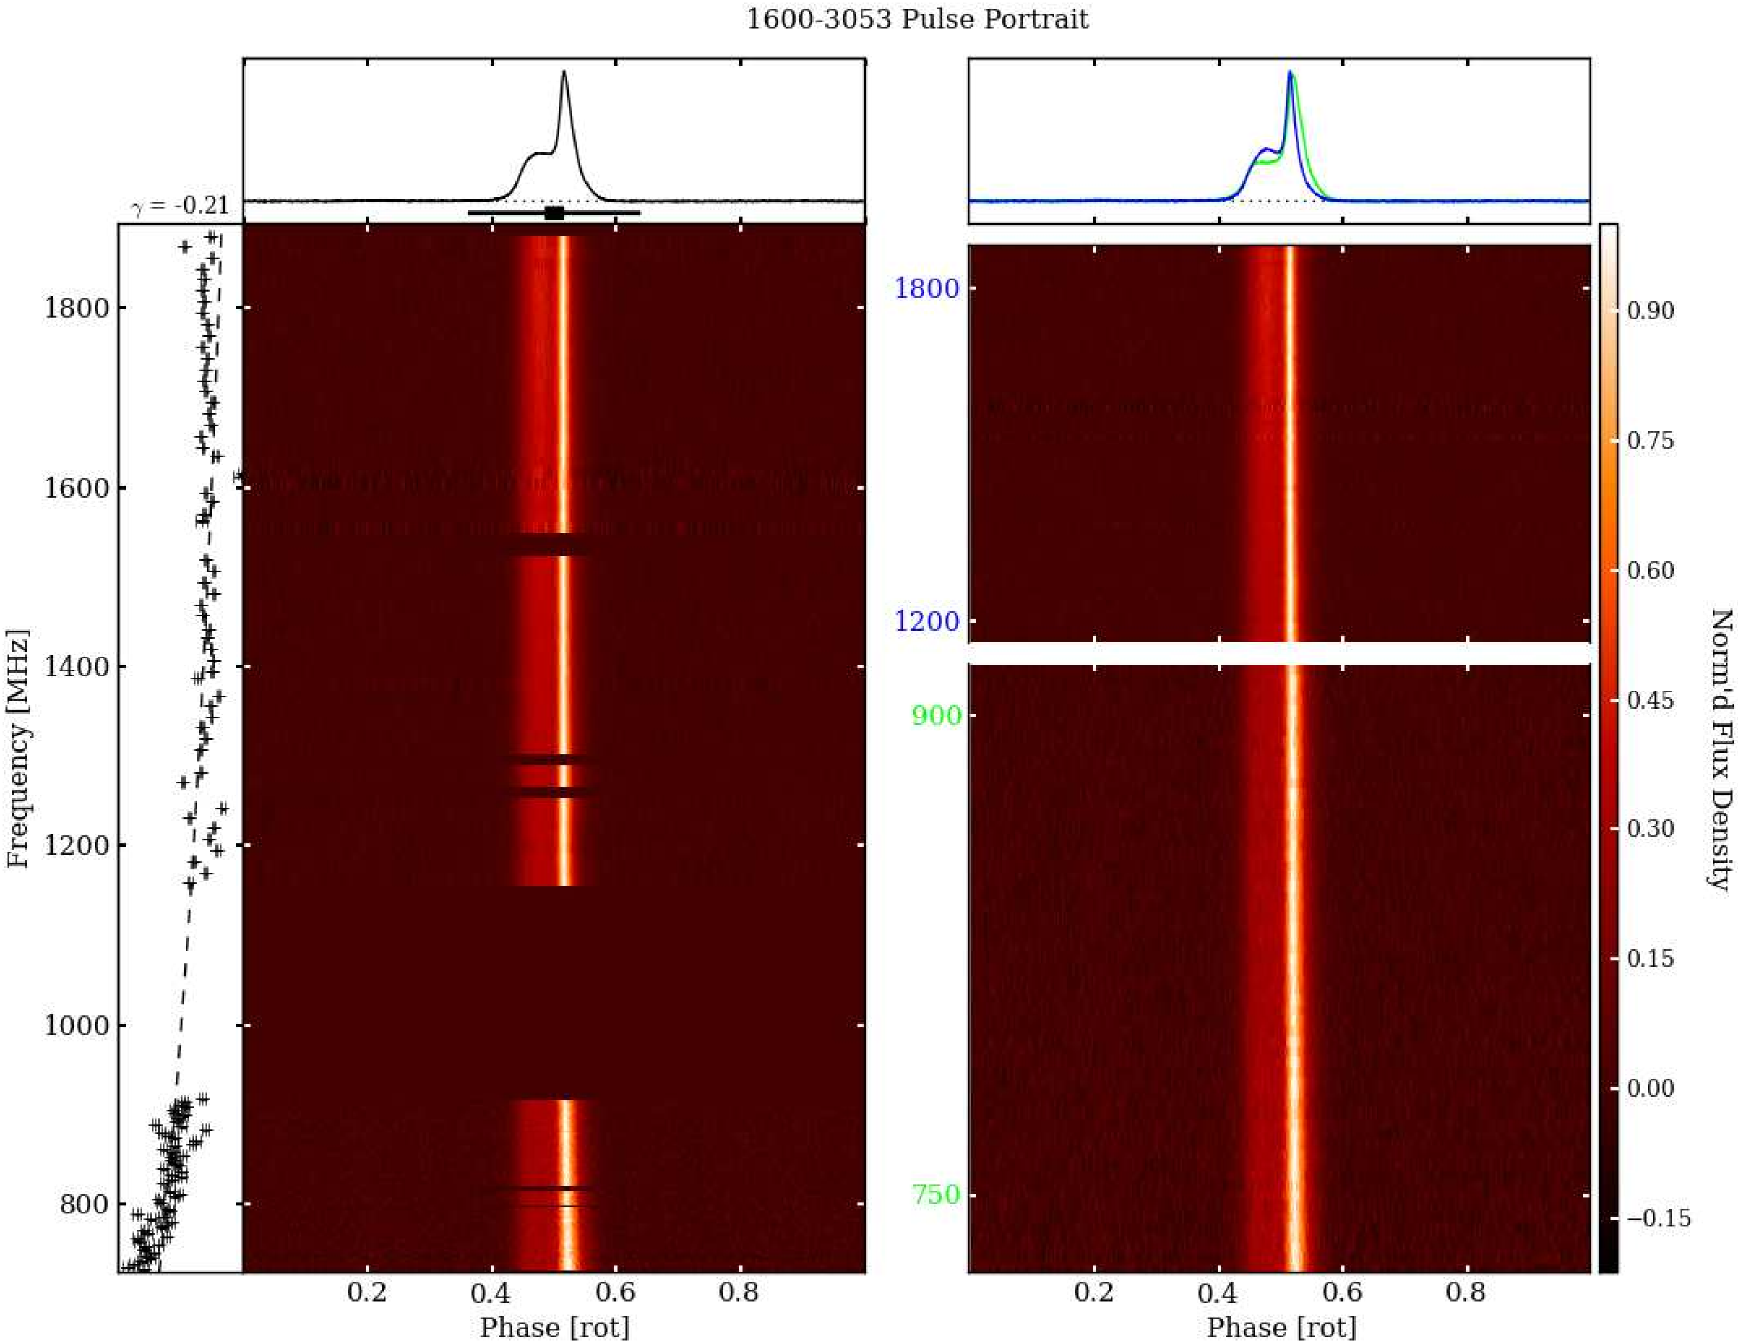

Fig. 4.174 – Concatenated portrait data and average profiles for J1600−3053.

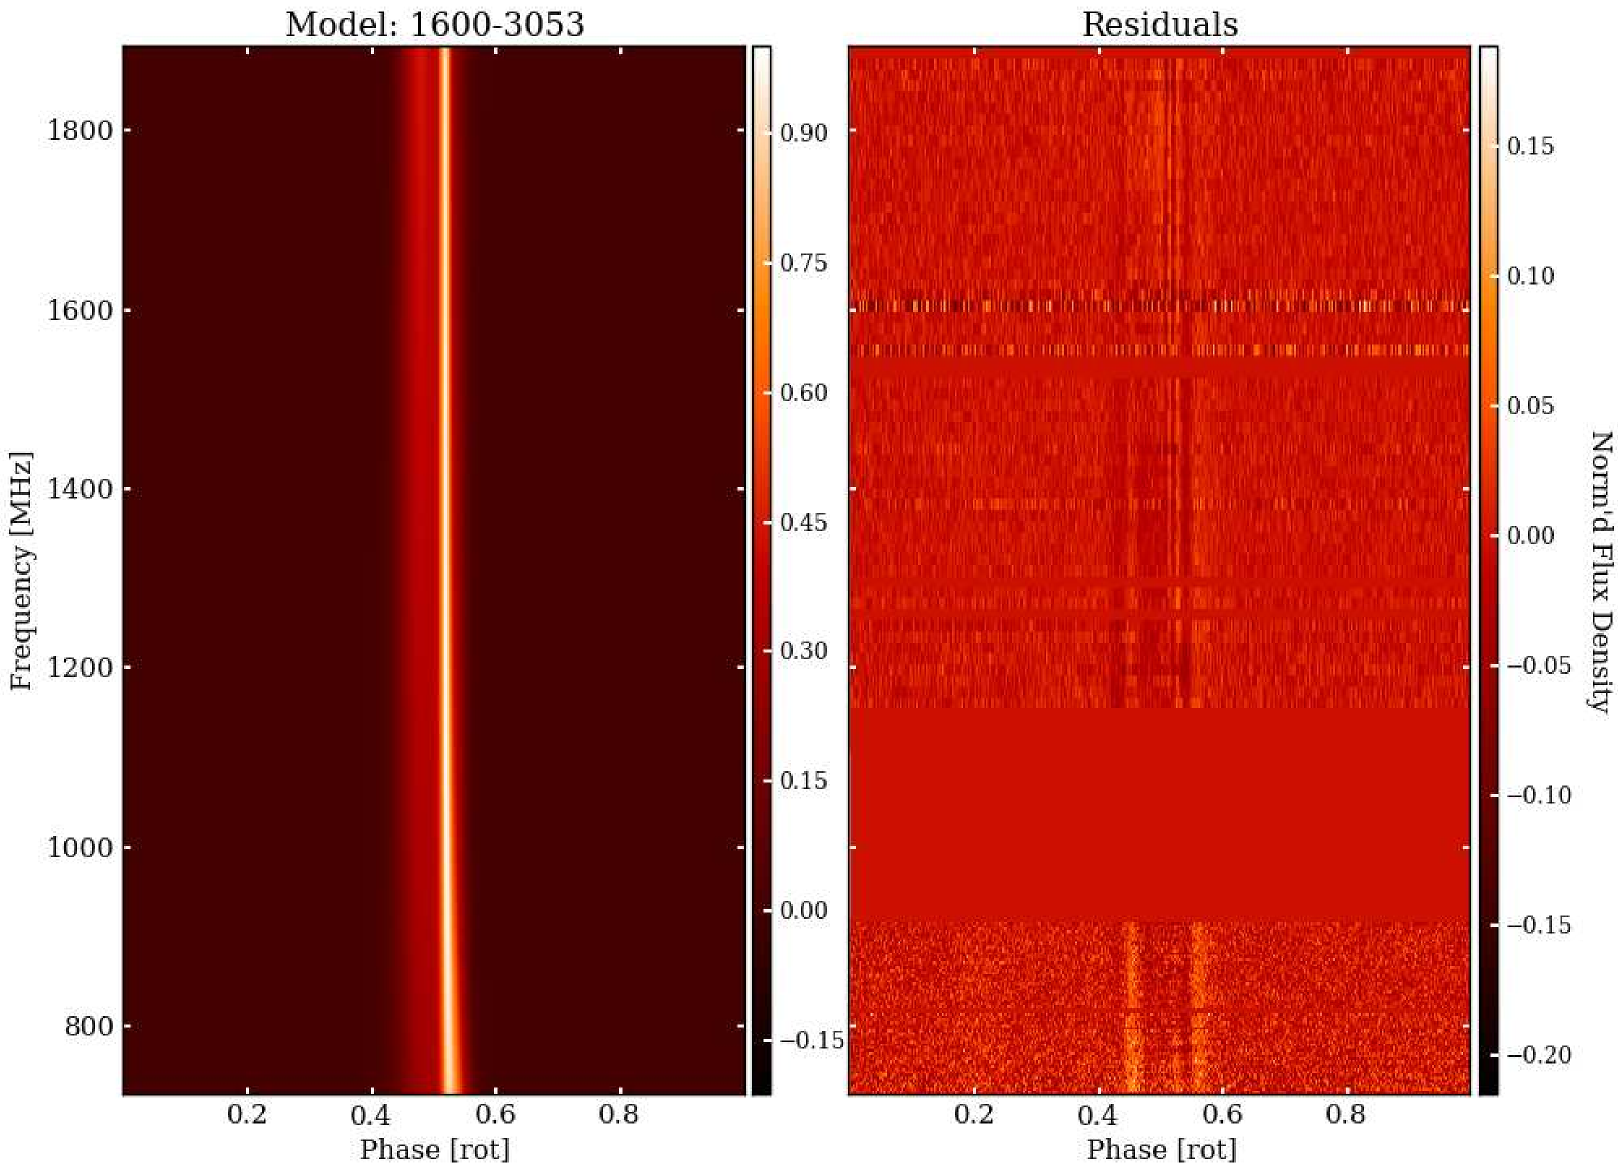

Fig. 4.175 – Fitted Gaussian component model and residuals for J1600−3053.

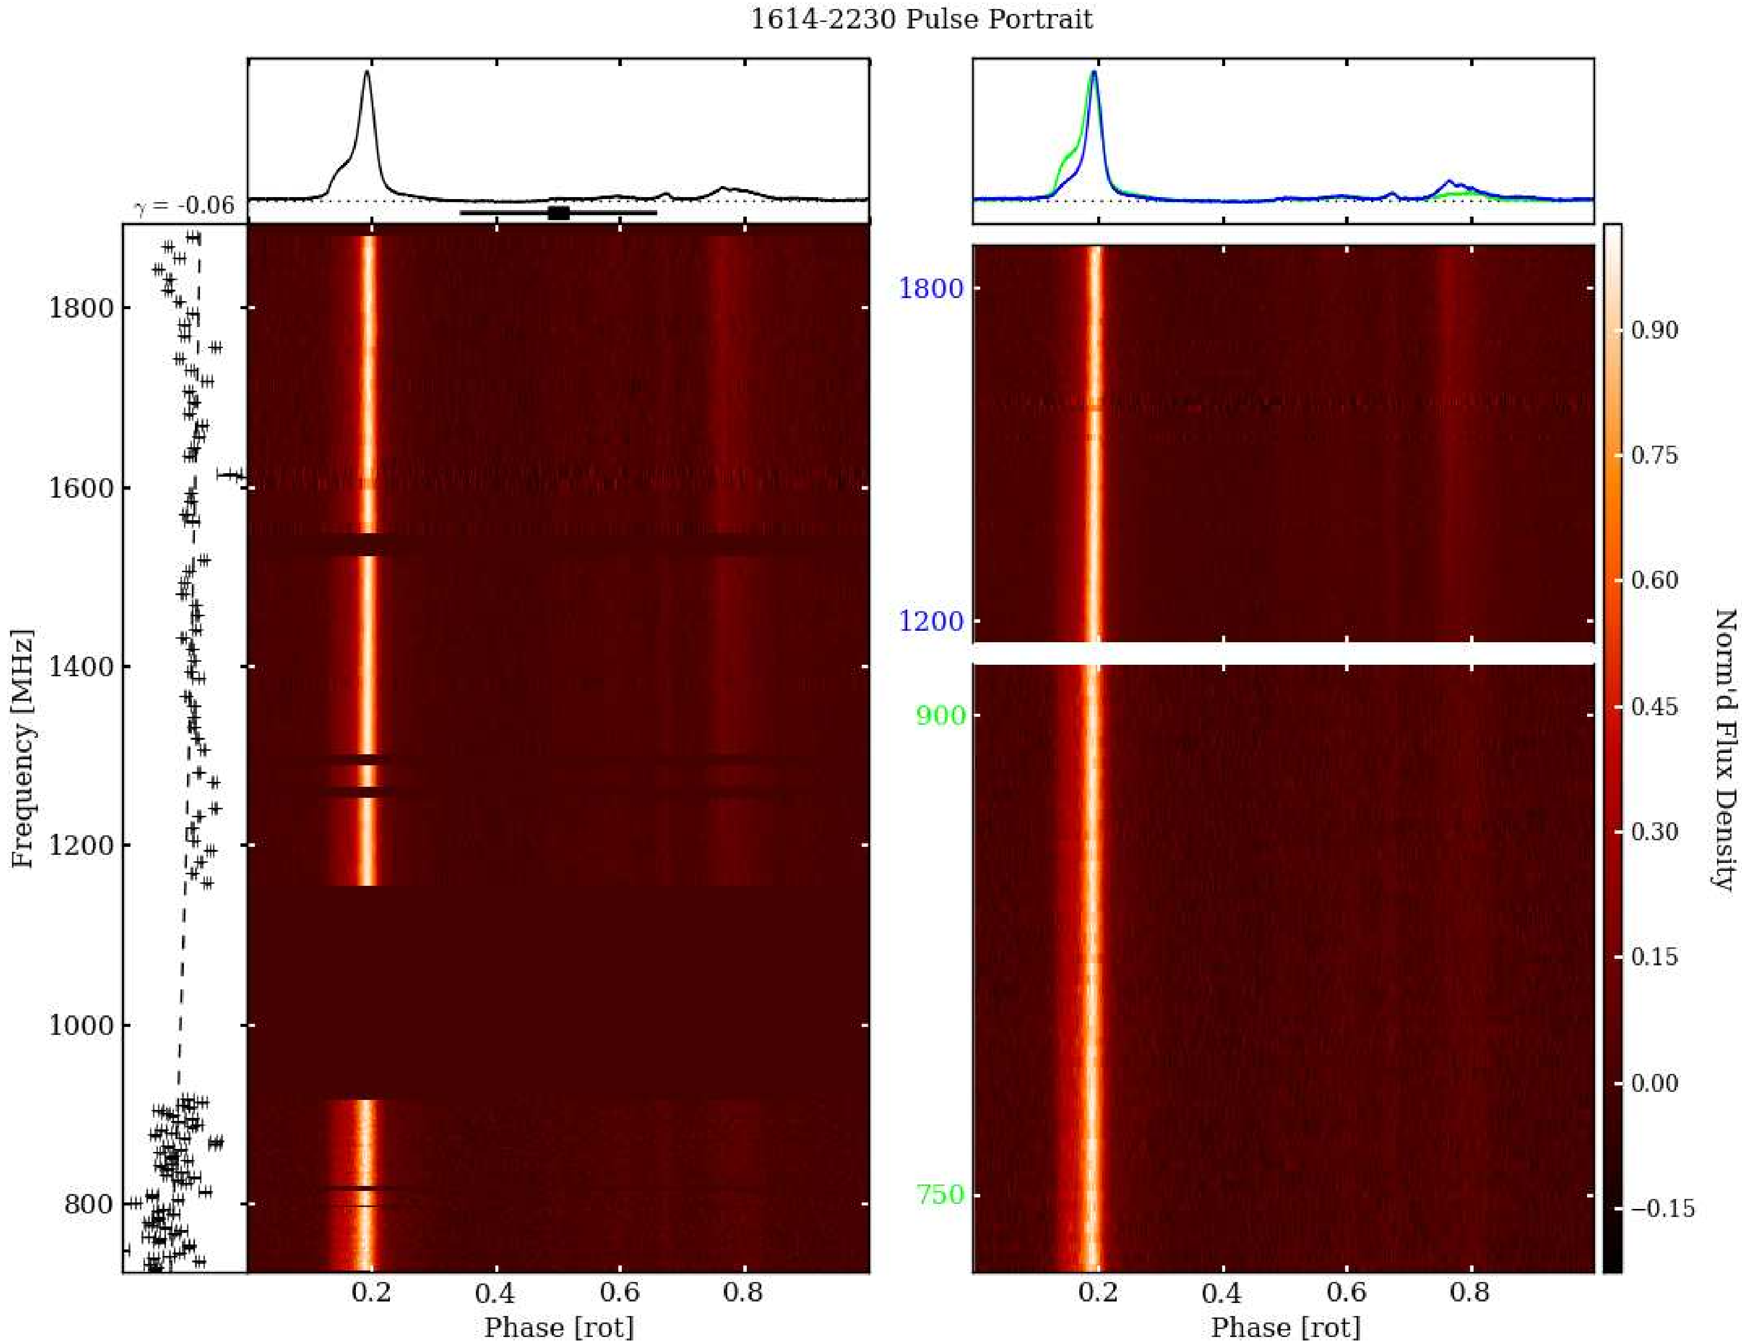

Fig. 4.176 – Concatenated portrait data and average profiles for J1614−2230.

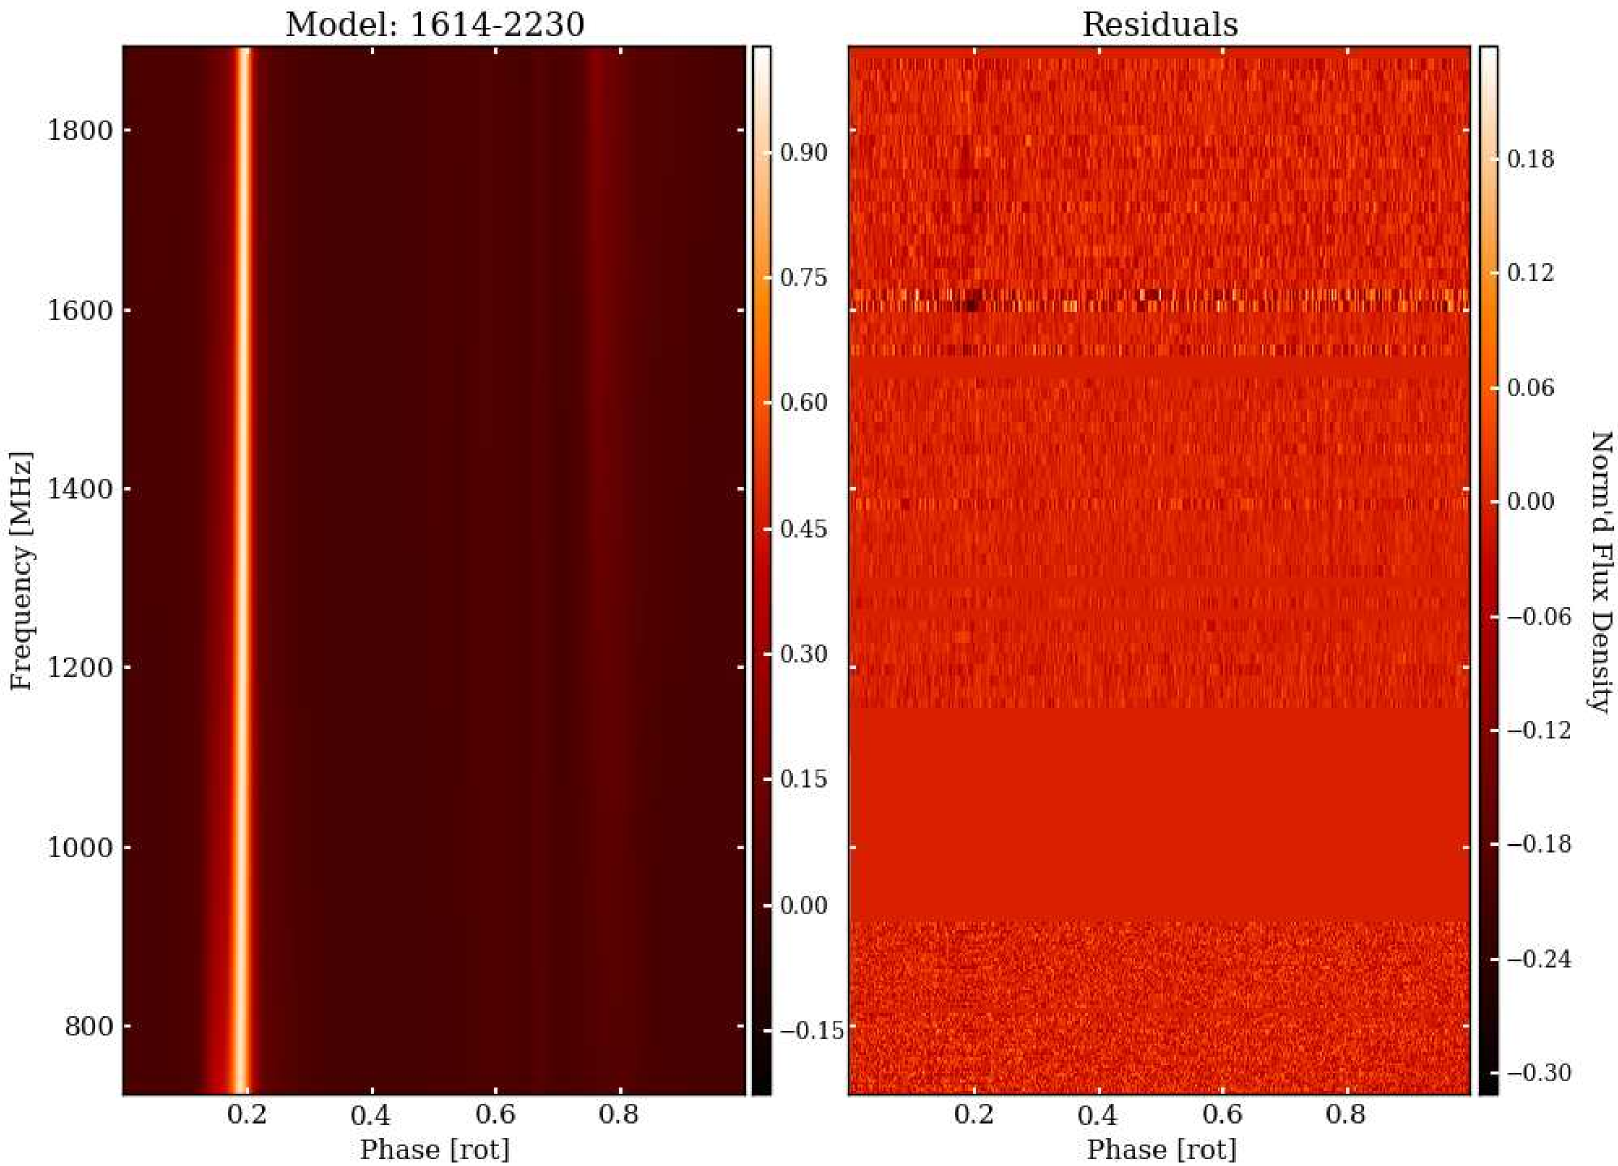

Fig. 4.177 – Fitted Gaussian component model and residuals for J1614−2230.

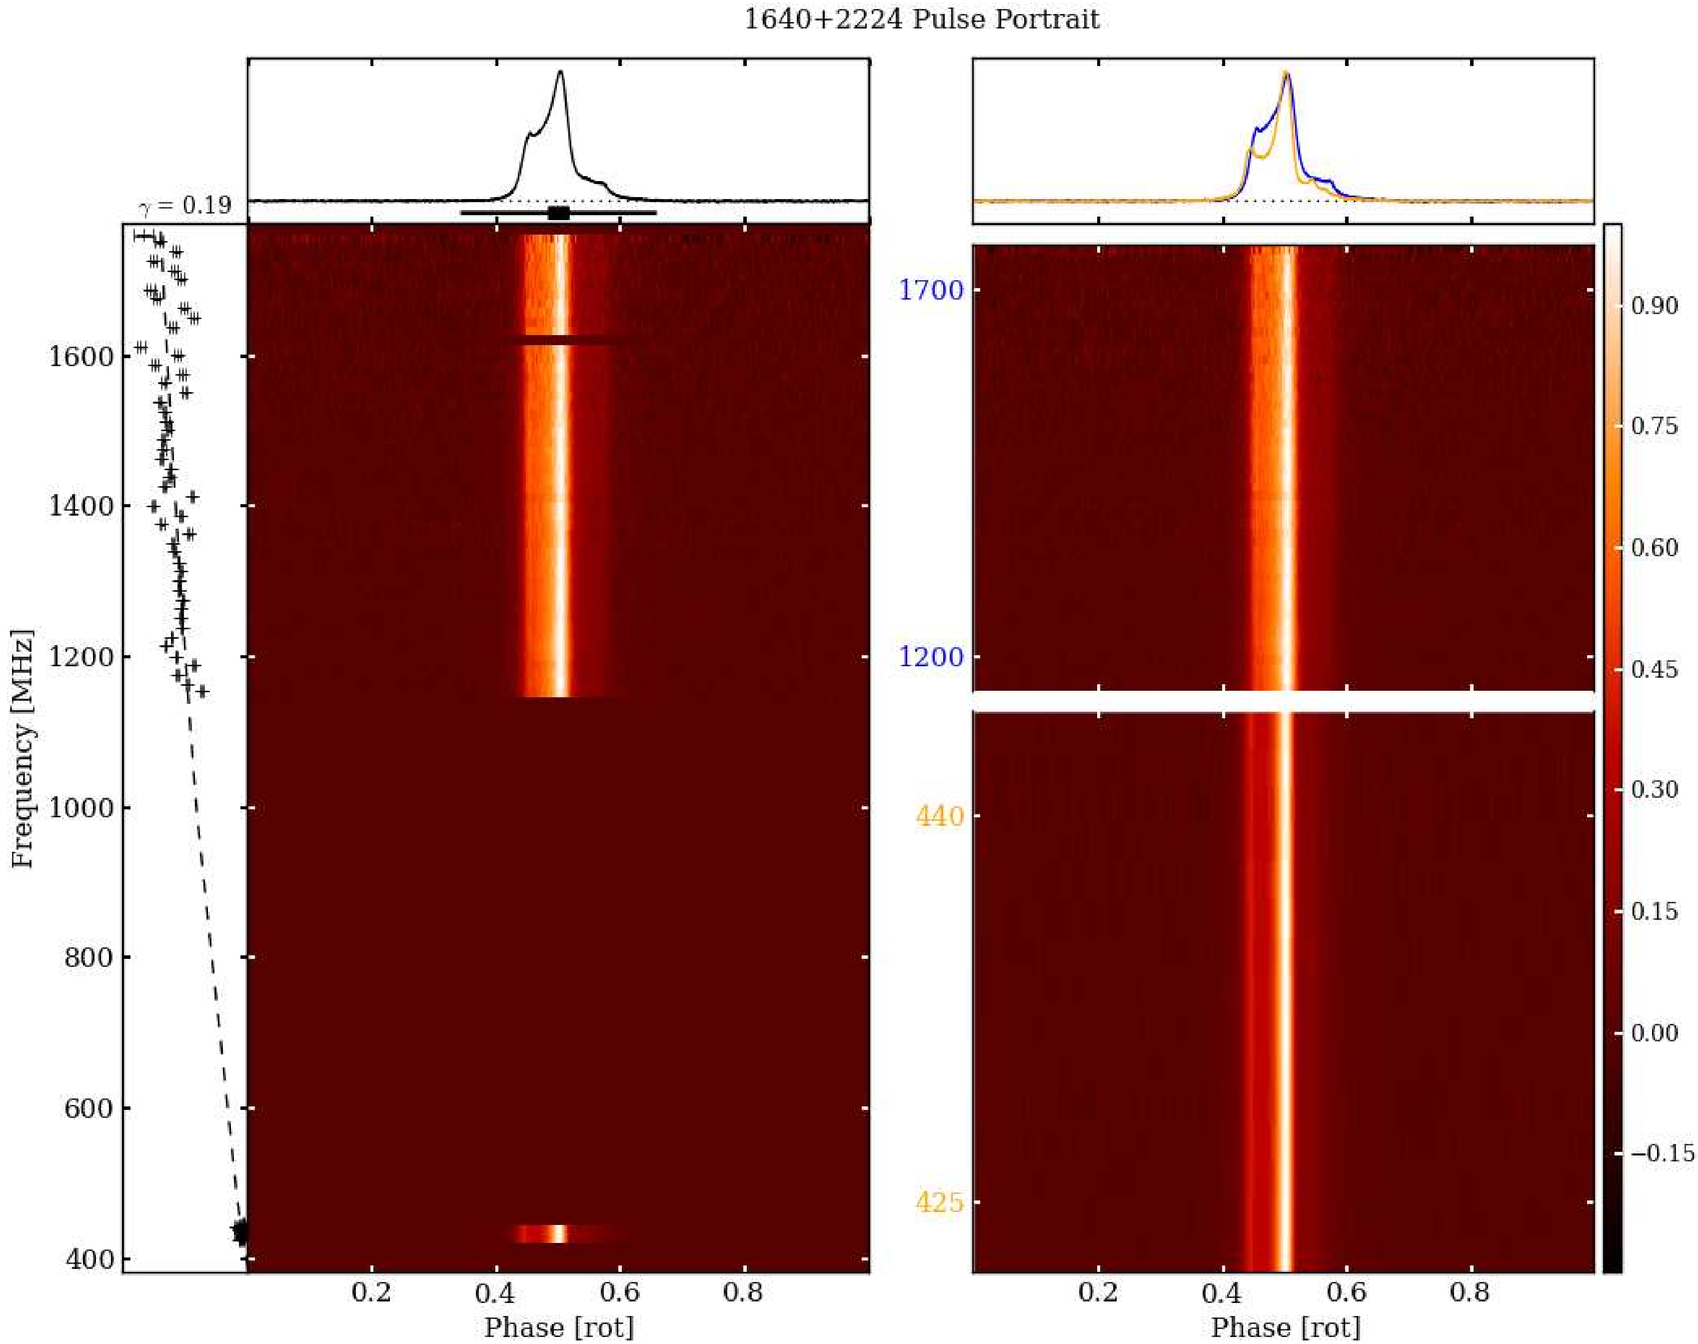

Fig. 4.178 – Concatenated portrait data and average profiles for J1640+2224.

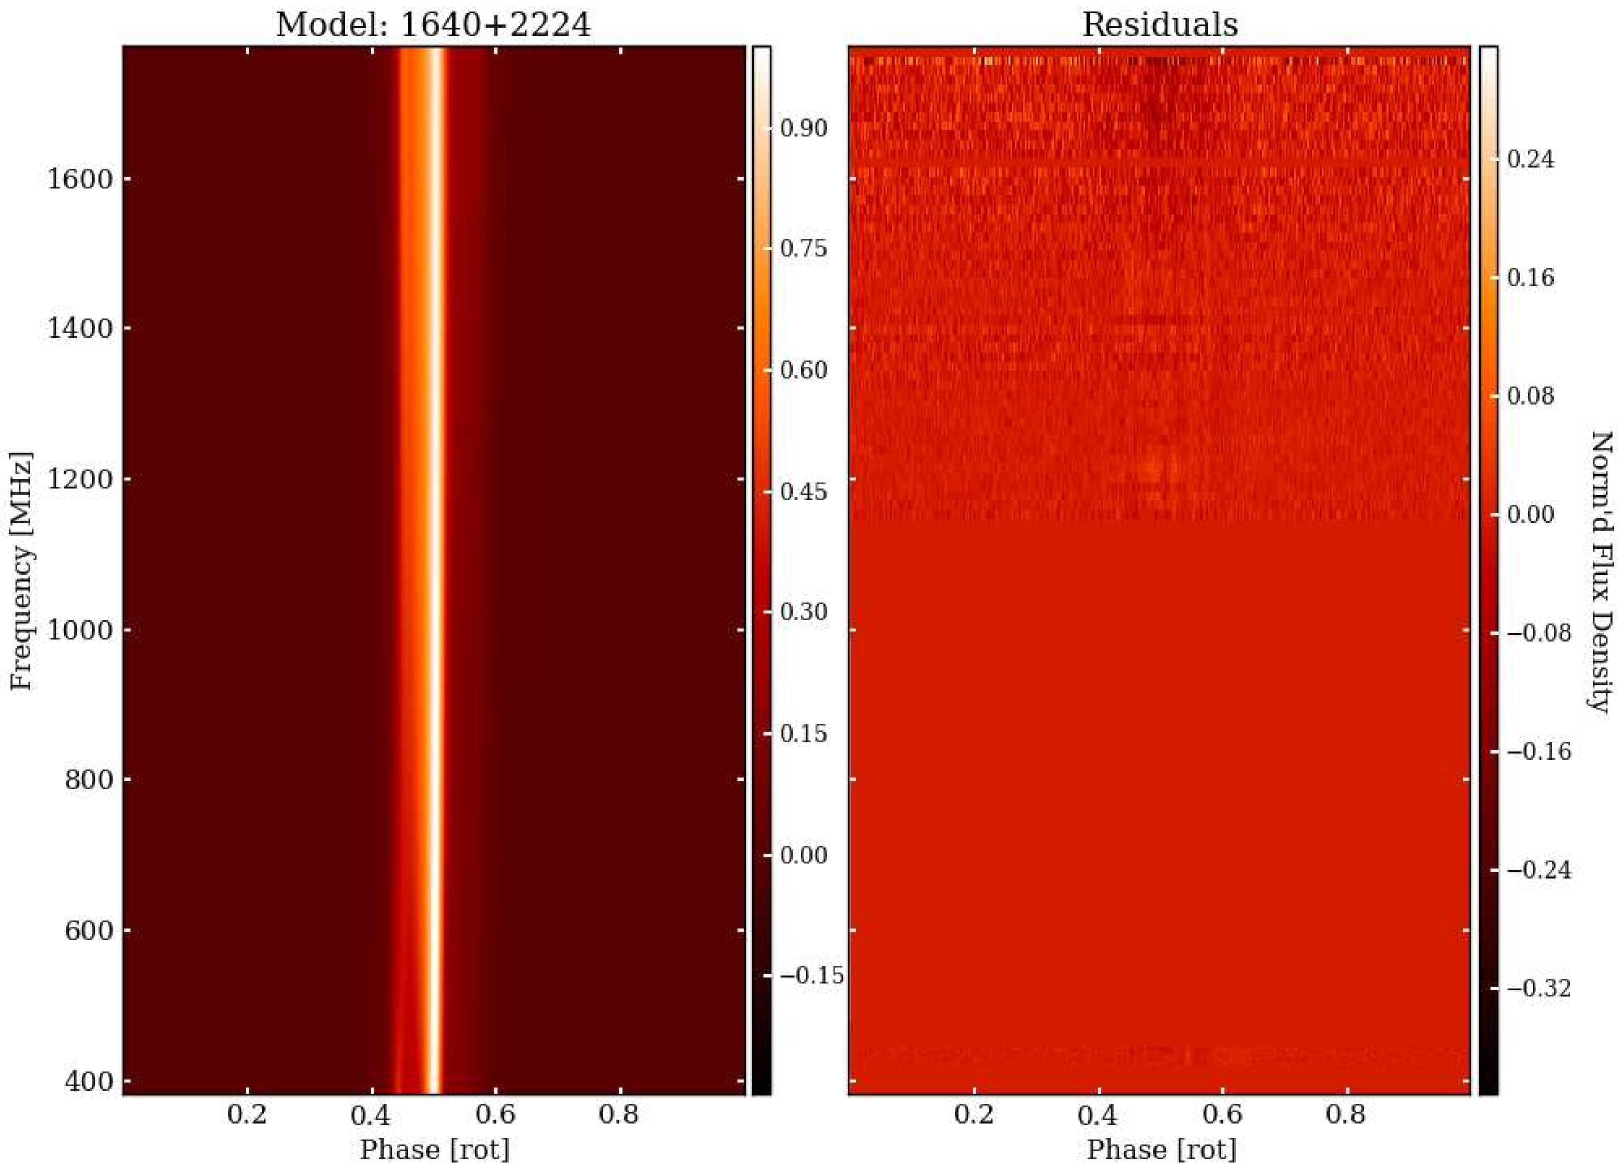

Fig. 4.179 – Fitted Gaussian component model and residuals for J1640+2224.

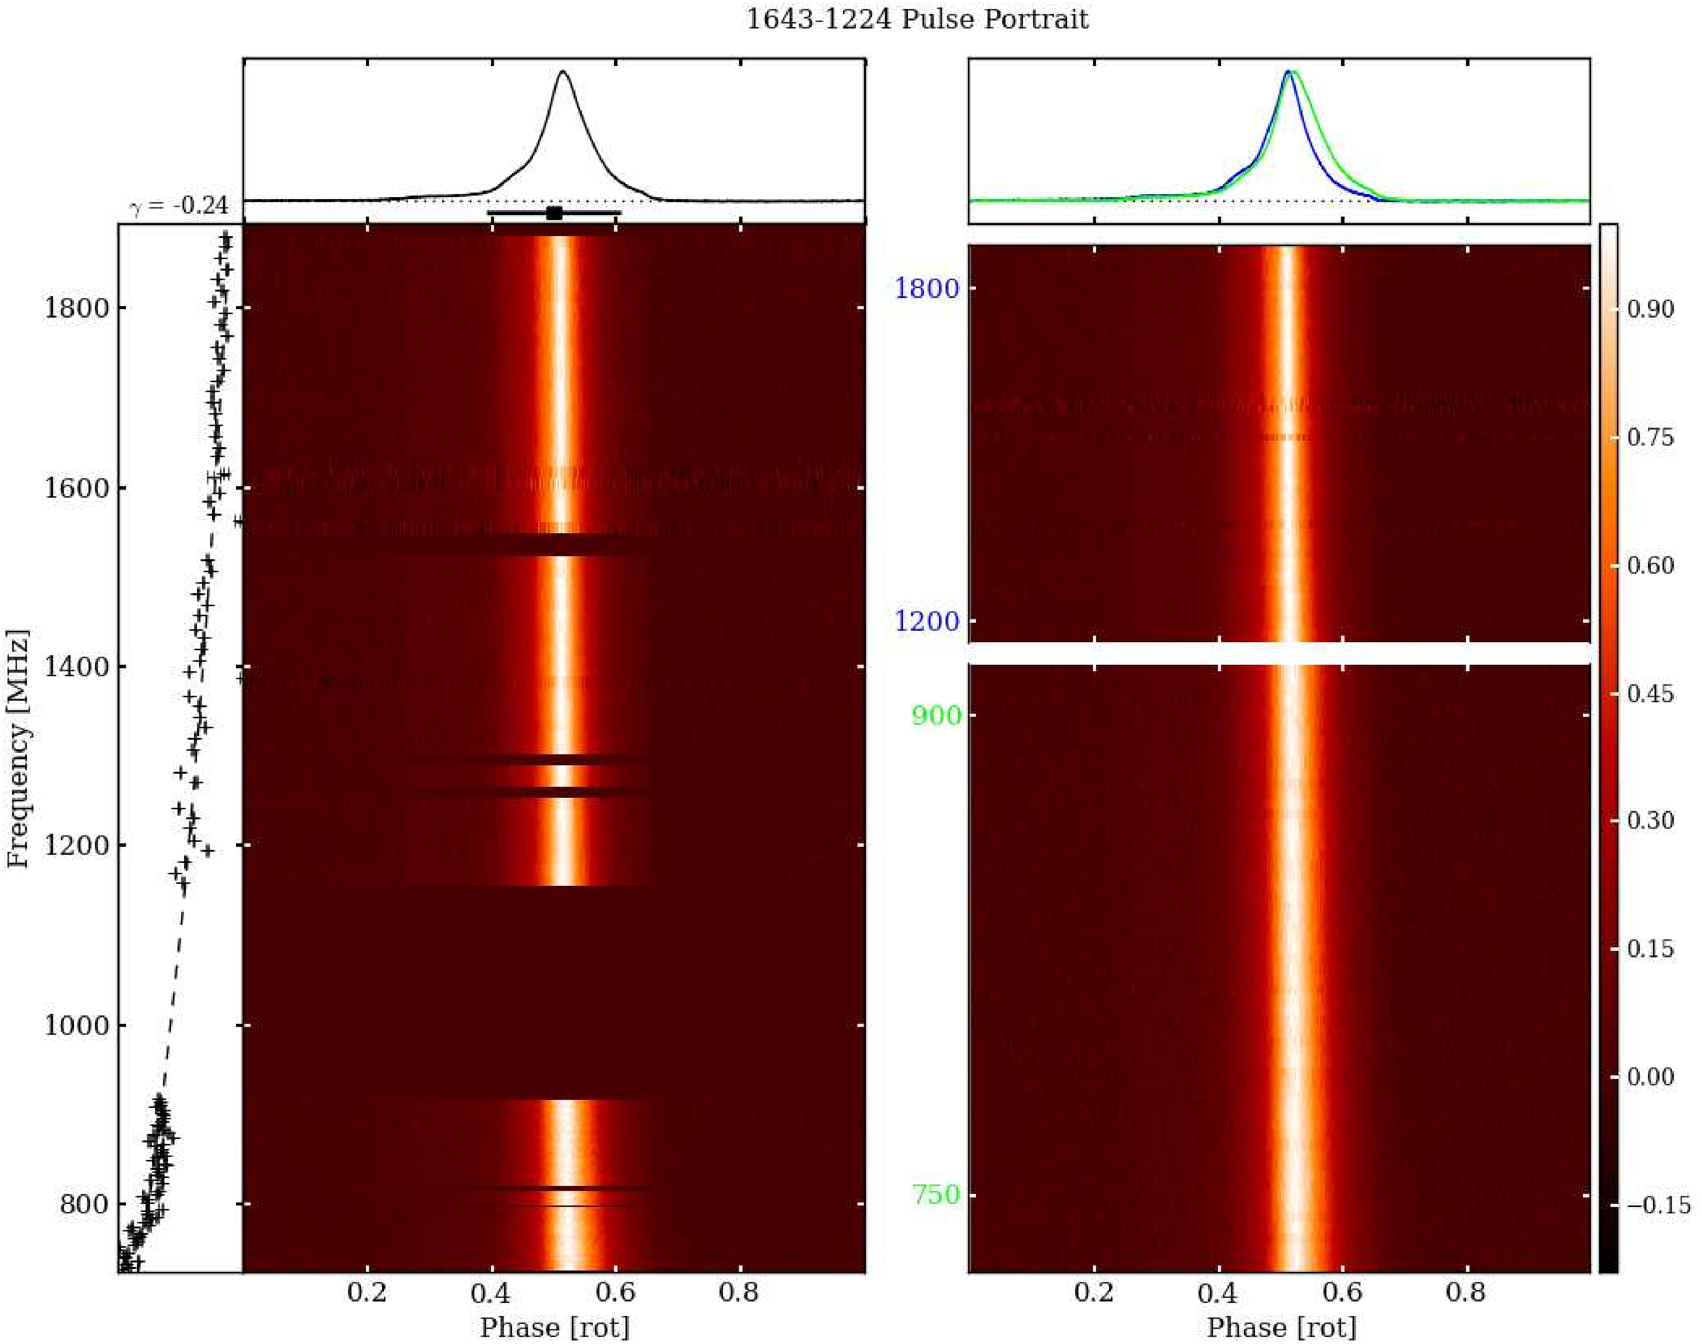

Fig. 4.180 – Concatenated portrait data and average profiles for J1643−1224.

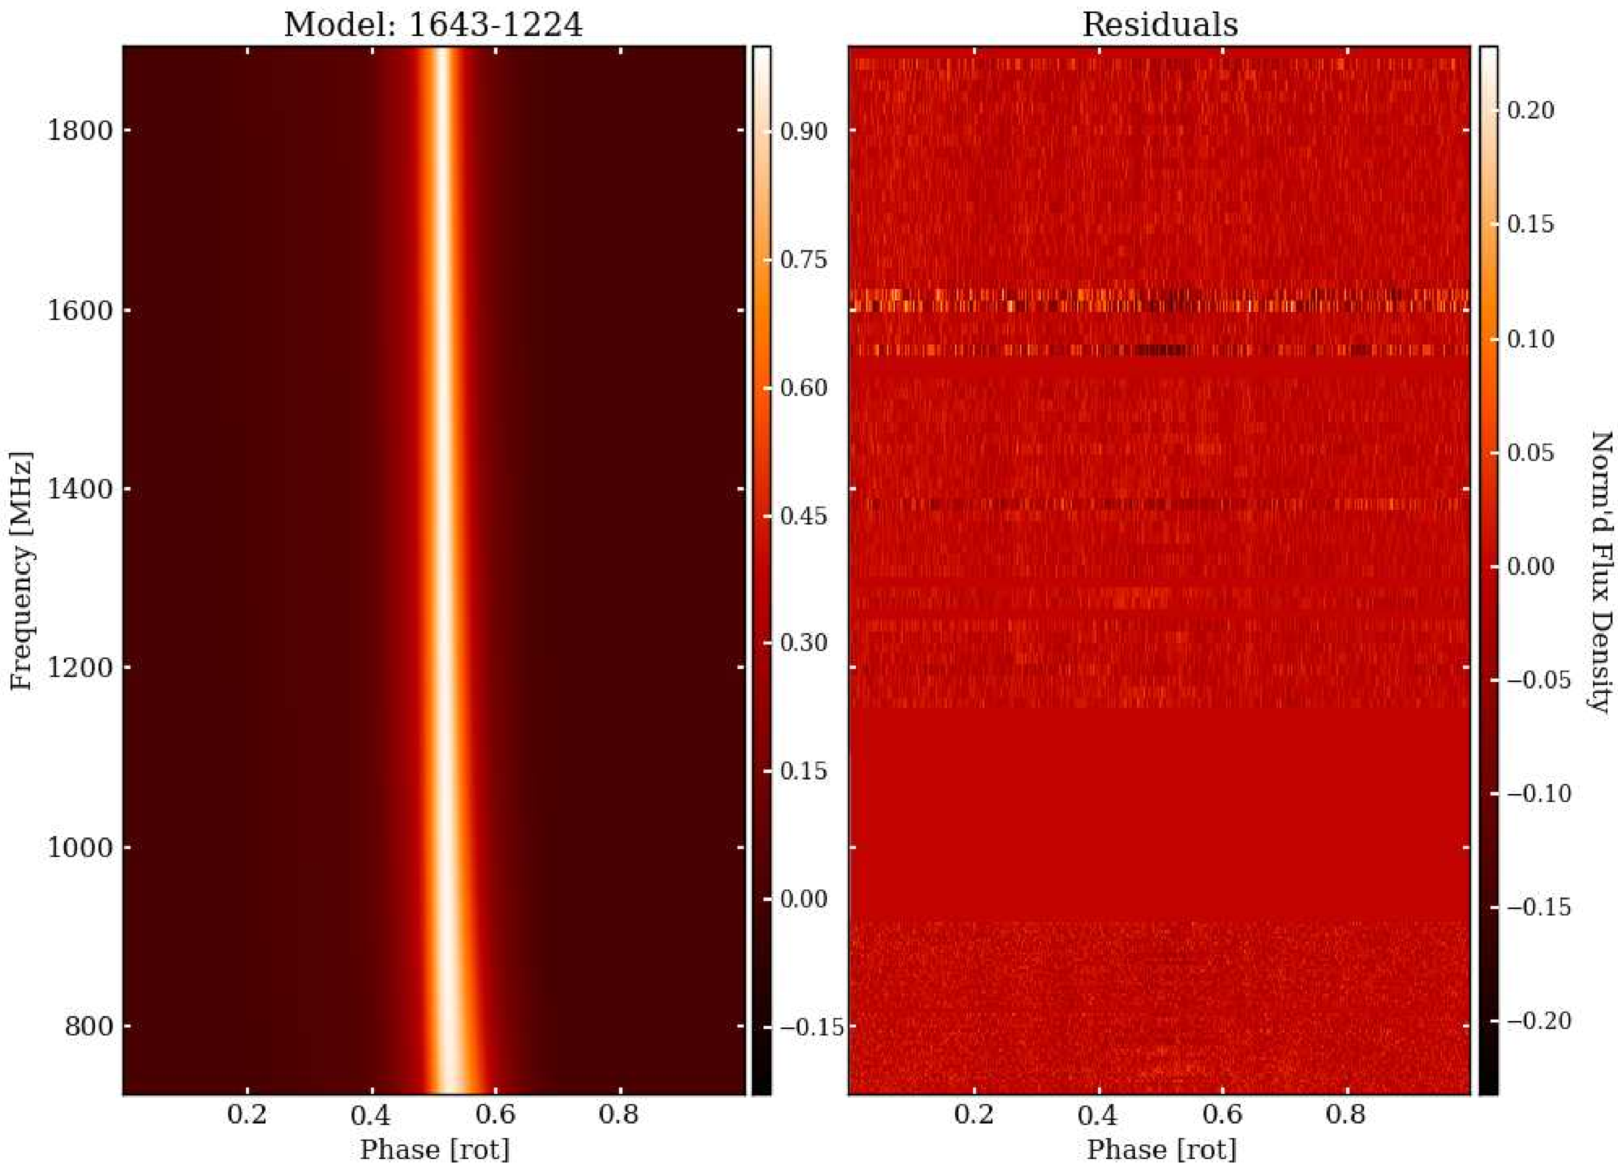

Fig. 4.181 – Fitted Gaussian component model and residuals for J1643−1224.

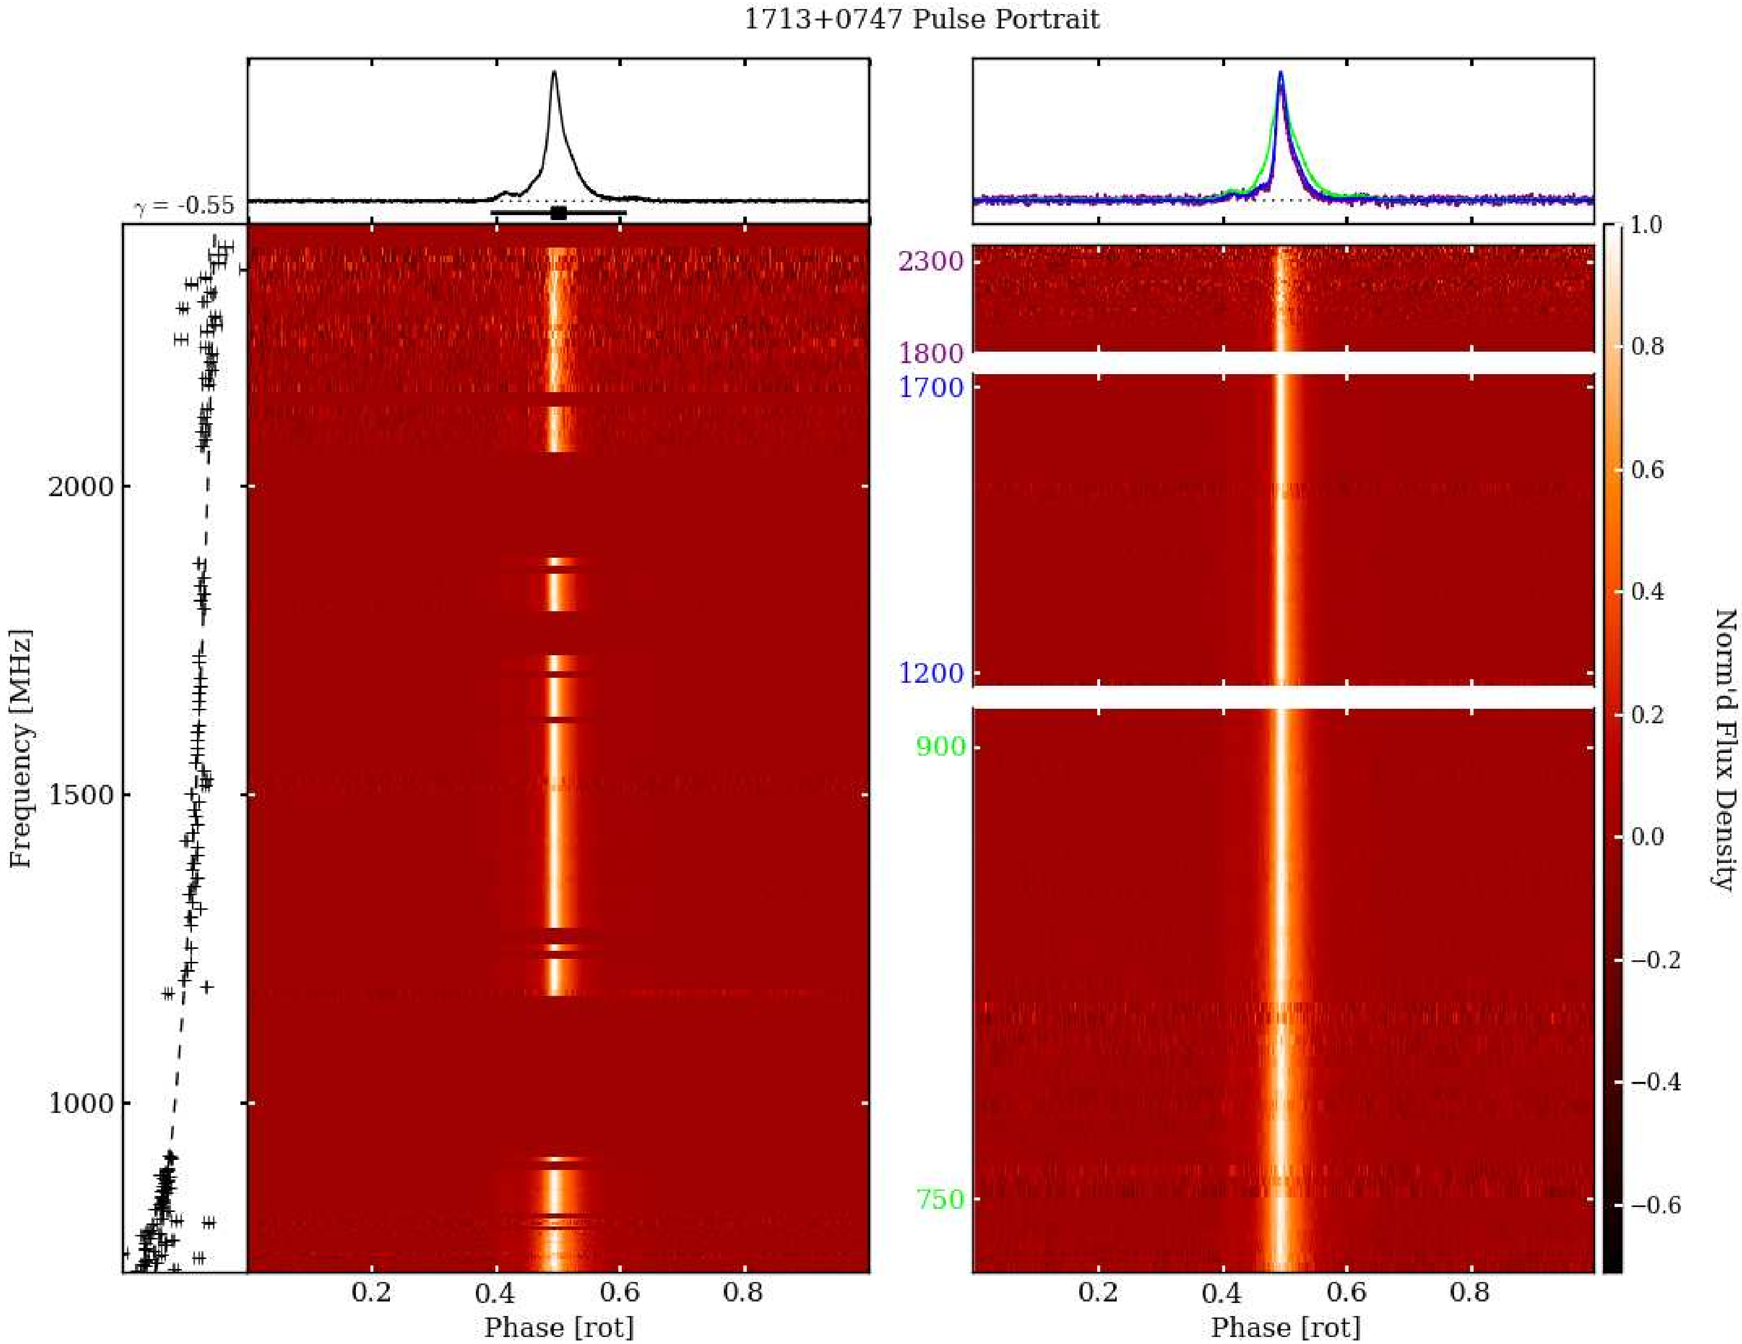

Fig. 4.182 – Concatenated portrait data and average profiles for J1713+0747.

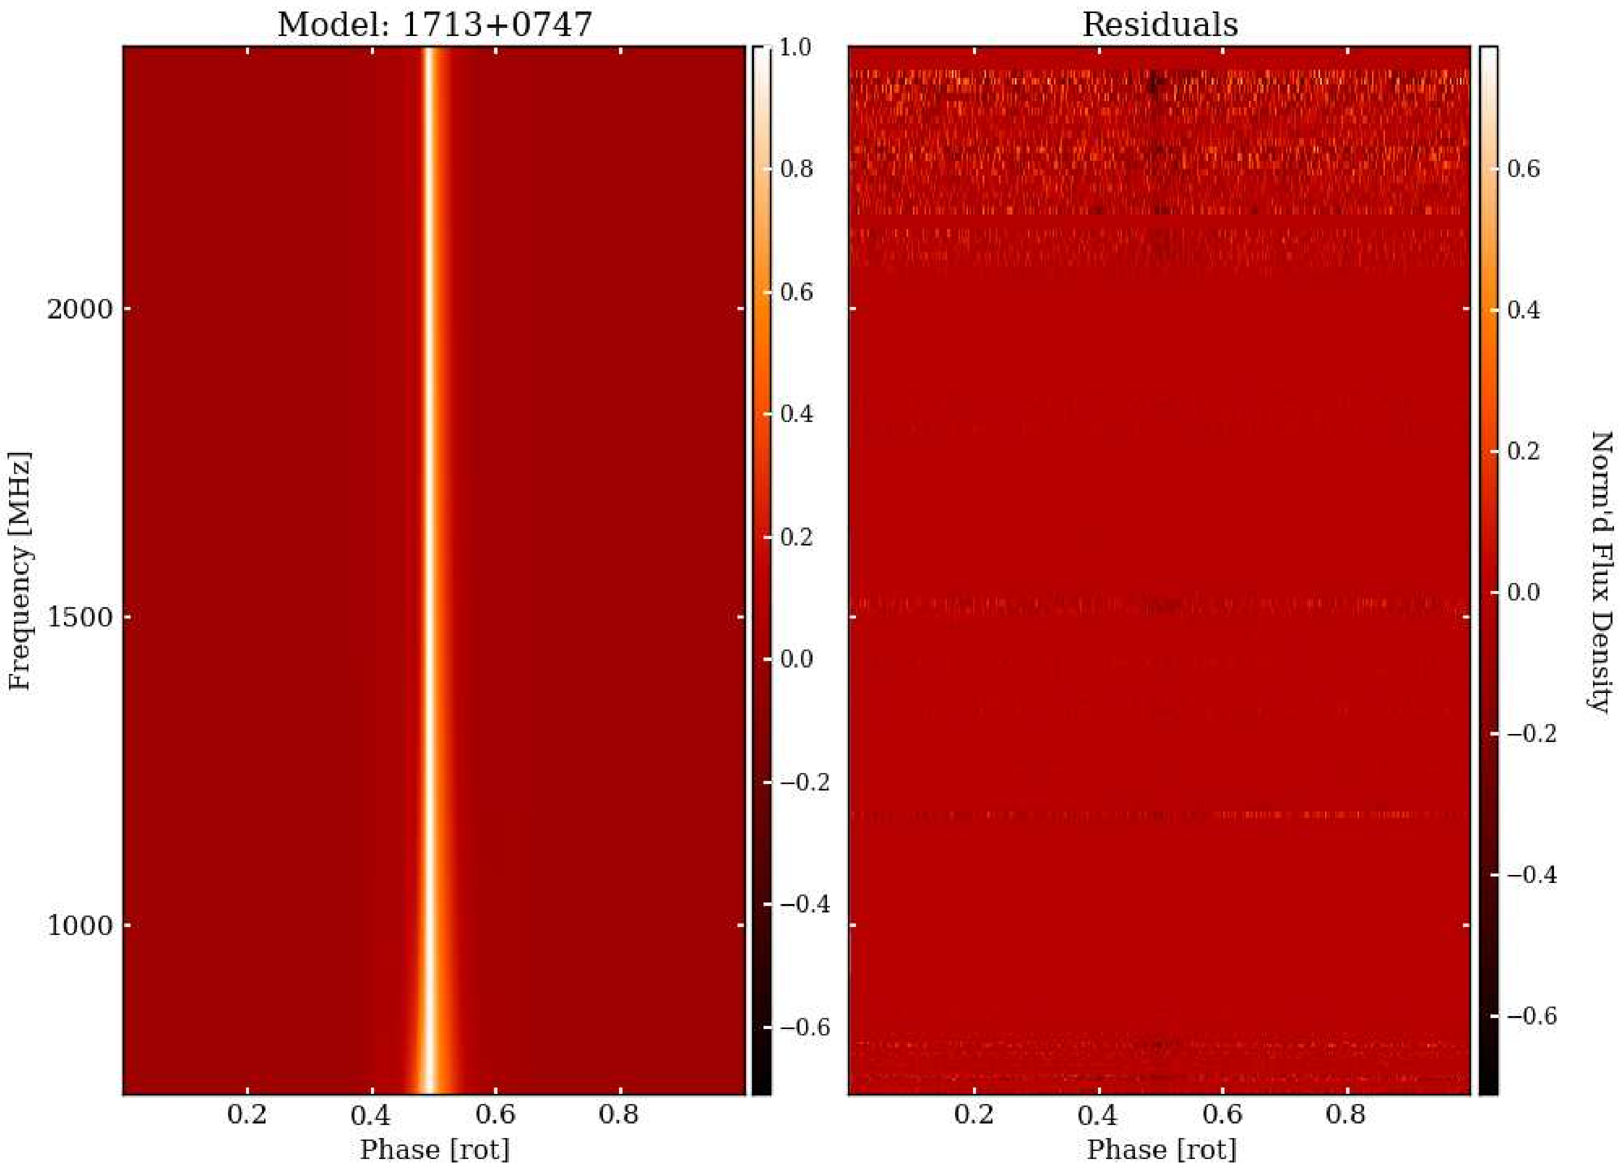

Fig. 4.183 – Fitted Gaussian component model and residuals for J1713+0747.

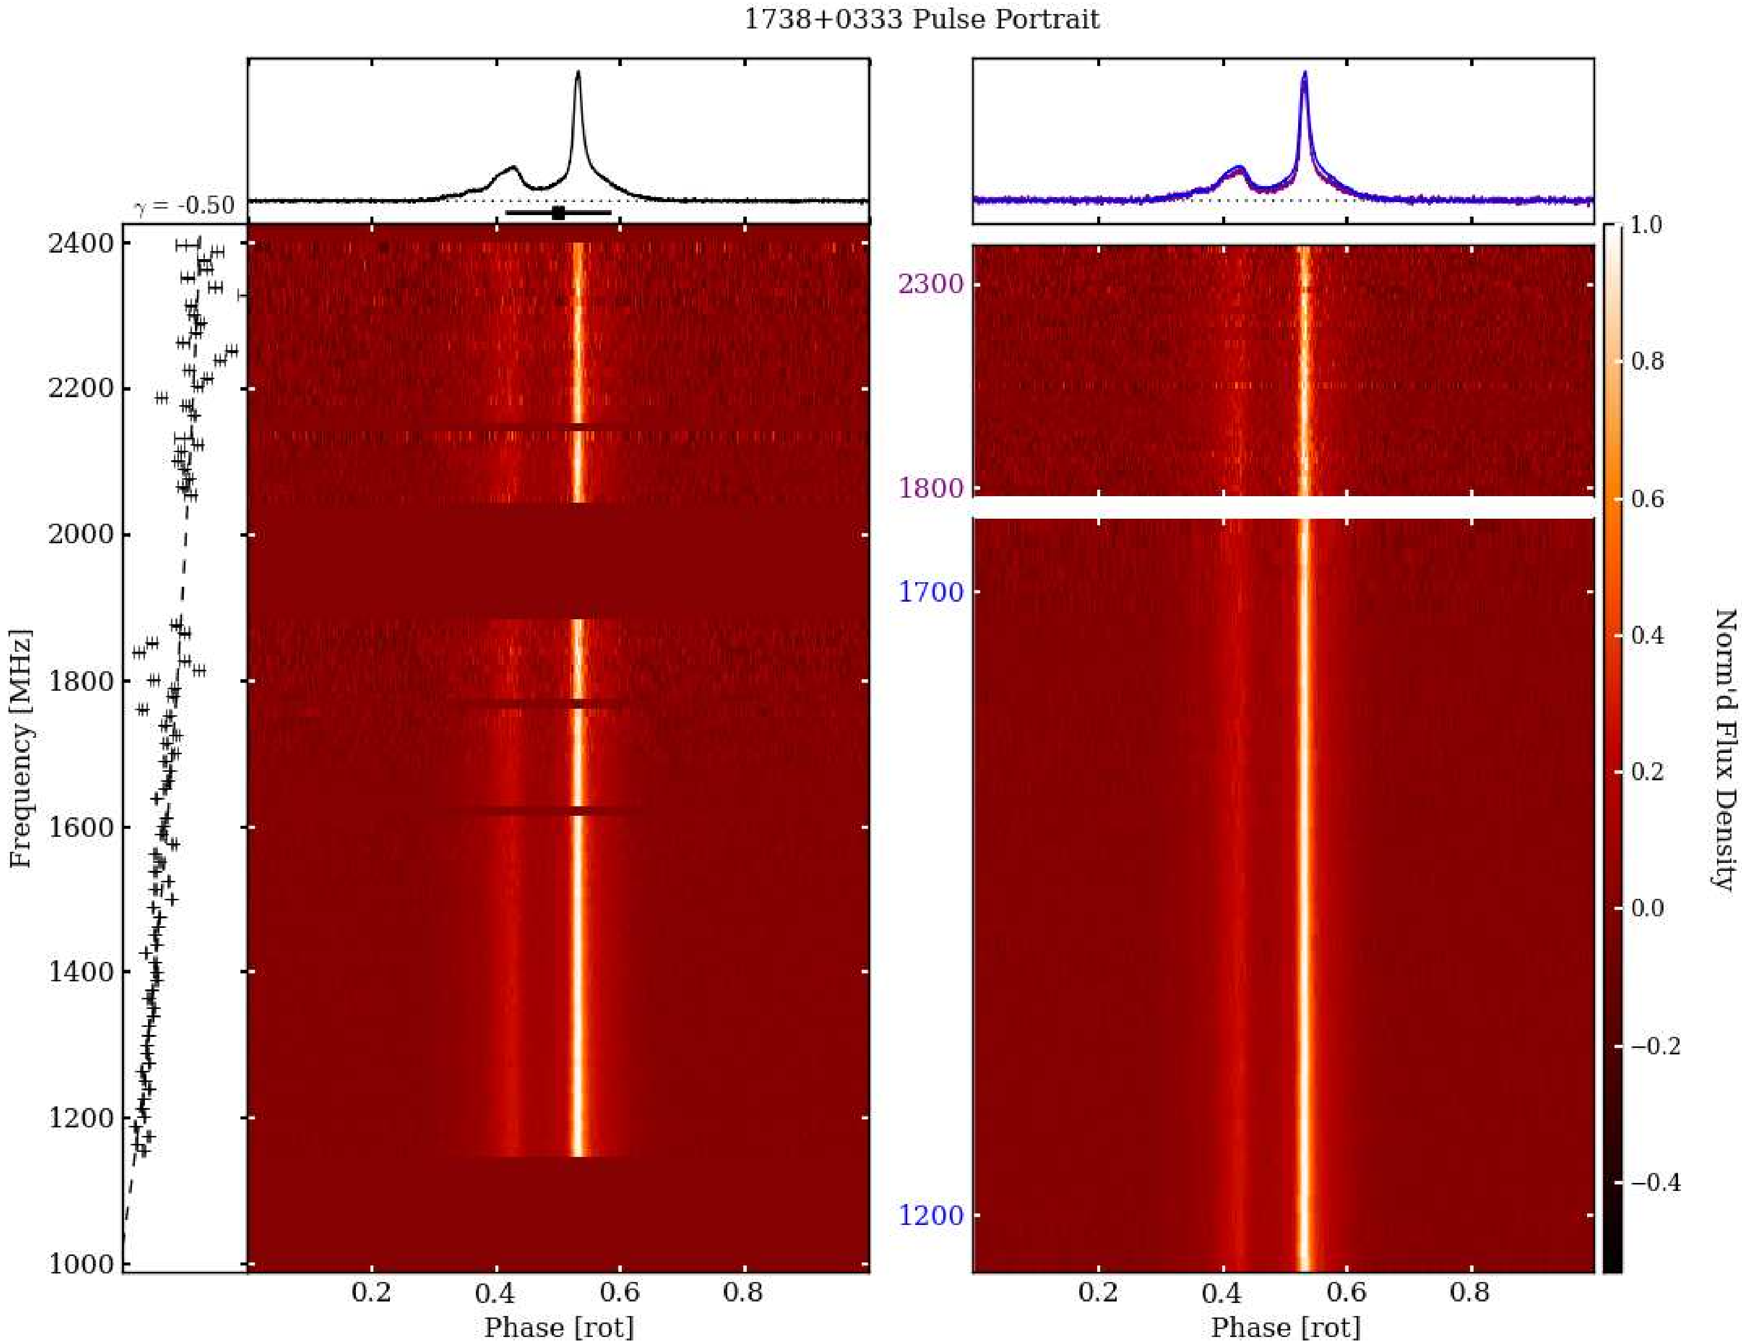

Fig. 4.184 – Concatenated portrait data and average profiles for J1738+0333.

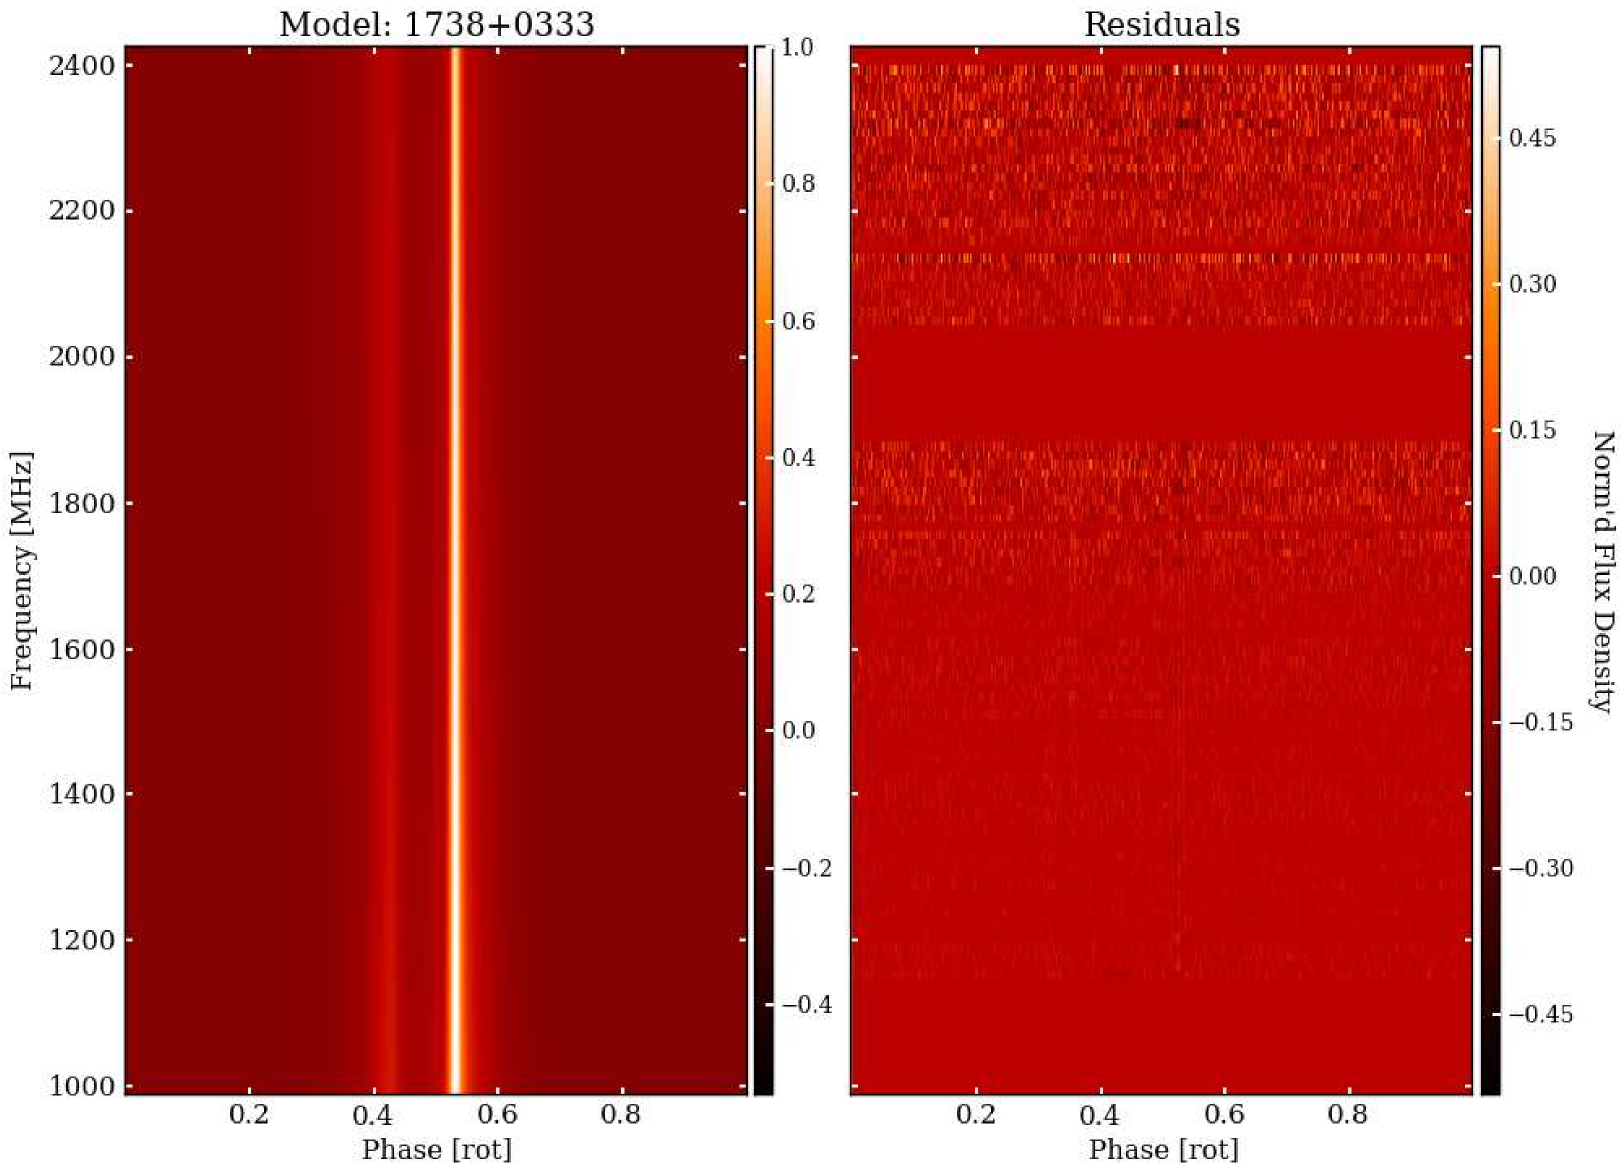

Fig. 4.185 – Fitted Gaussian component model and residuals for J1738+0333.

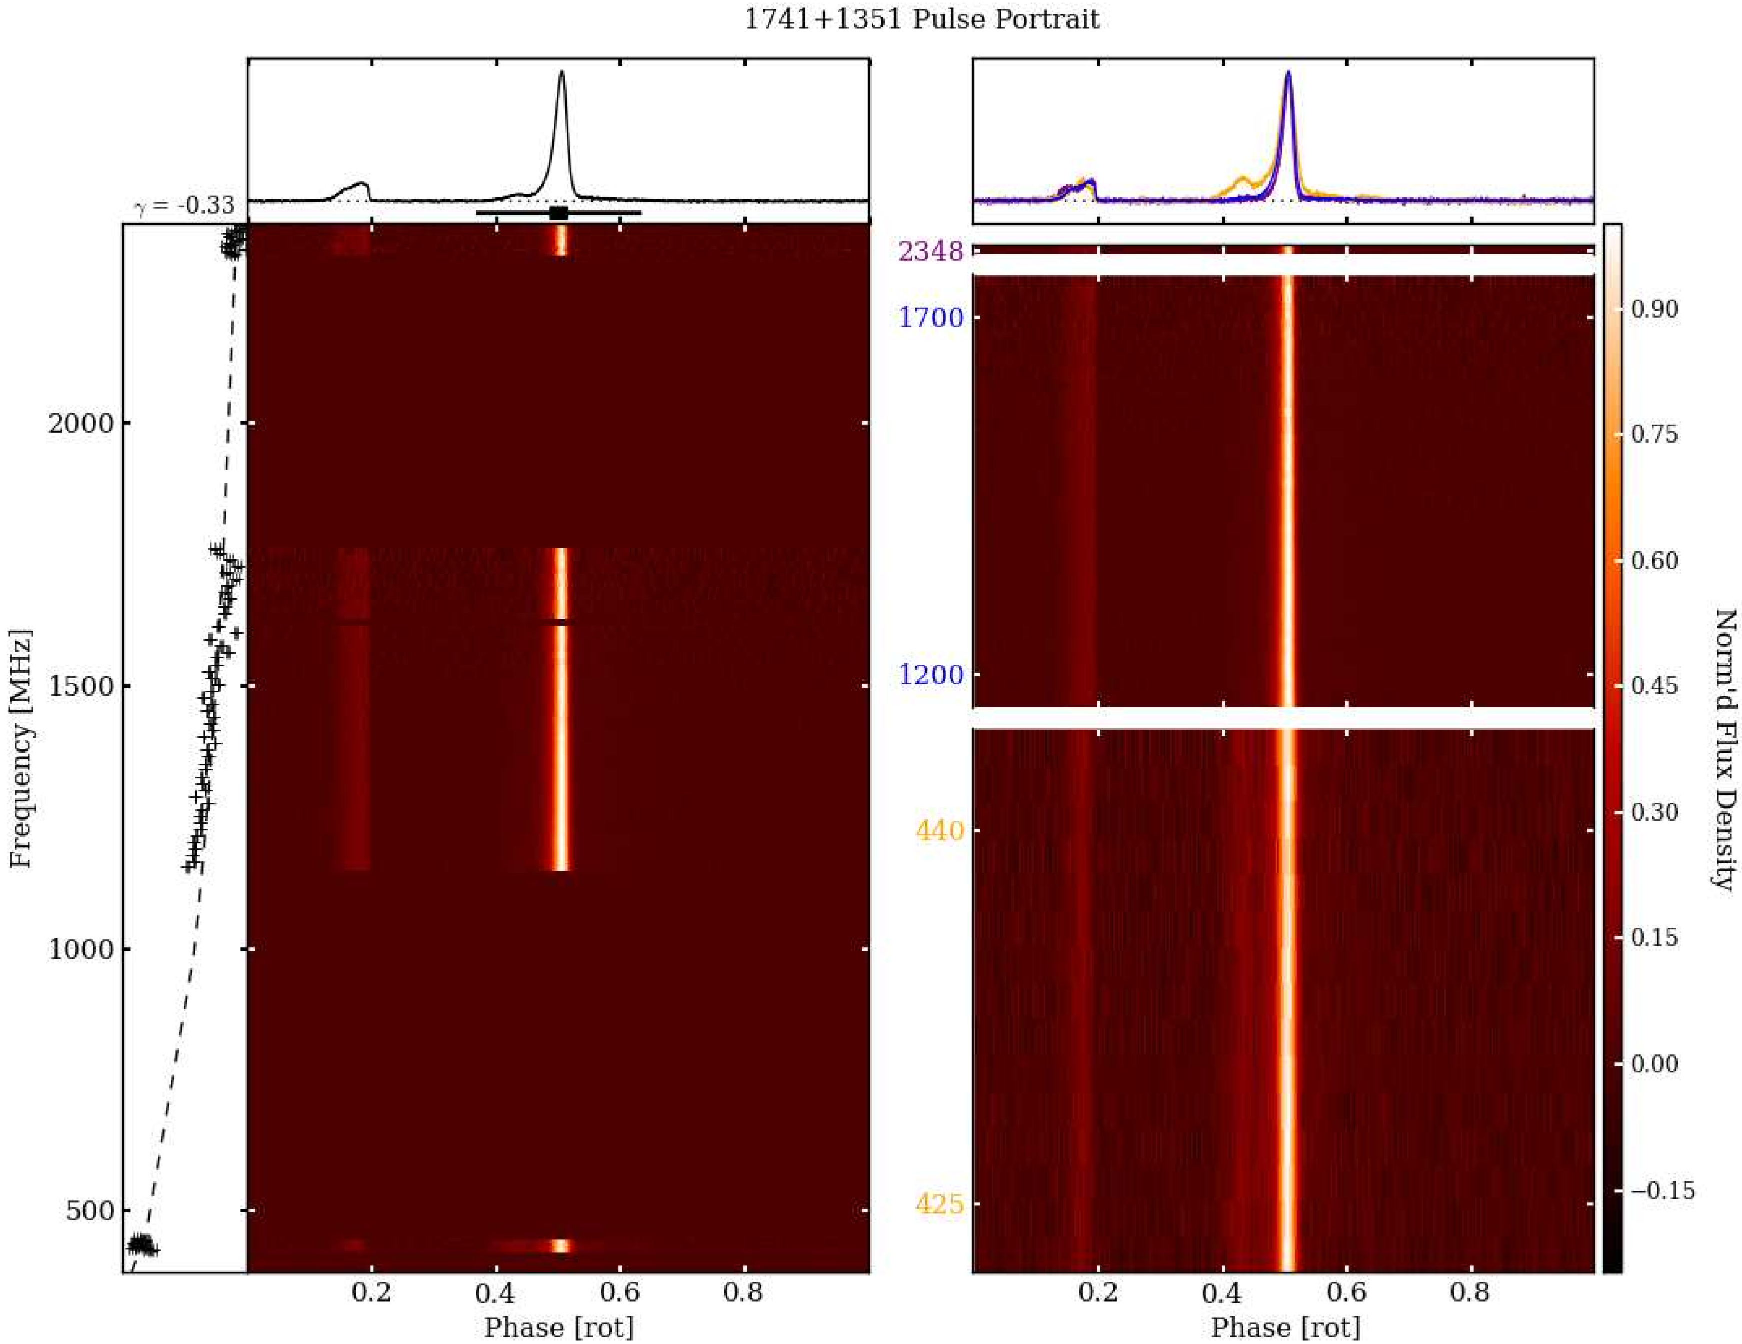

Fig. 4.186 – Concatenated portrait data and average profiles for J1741+1351.

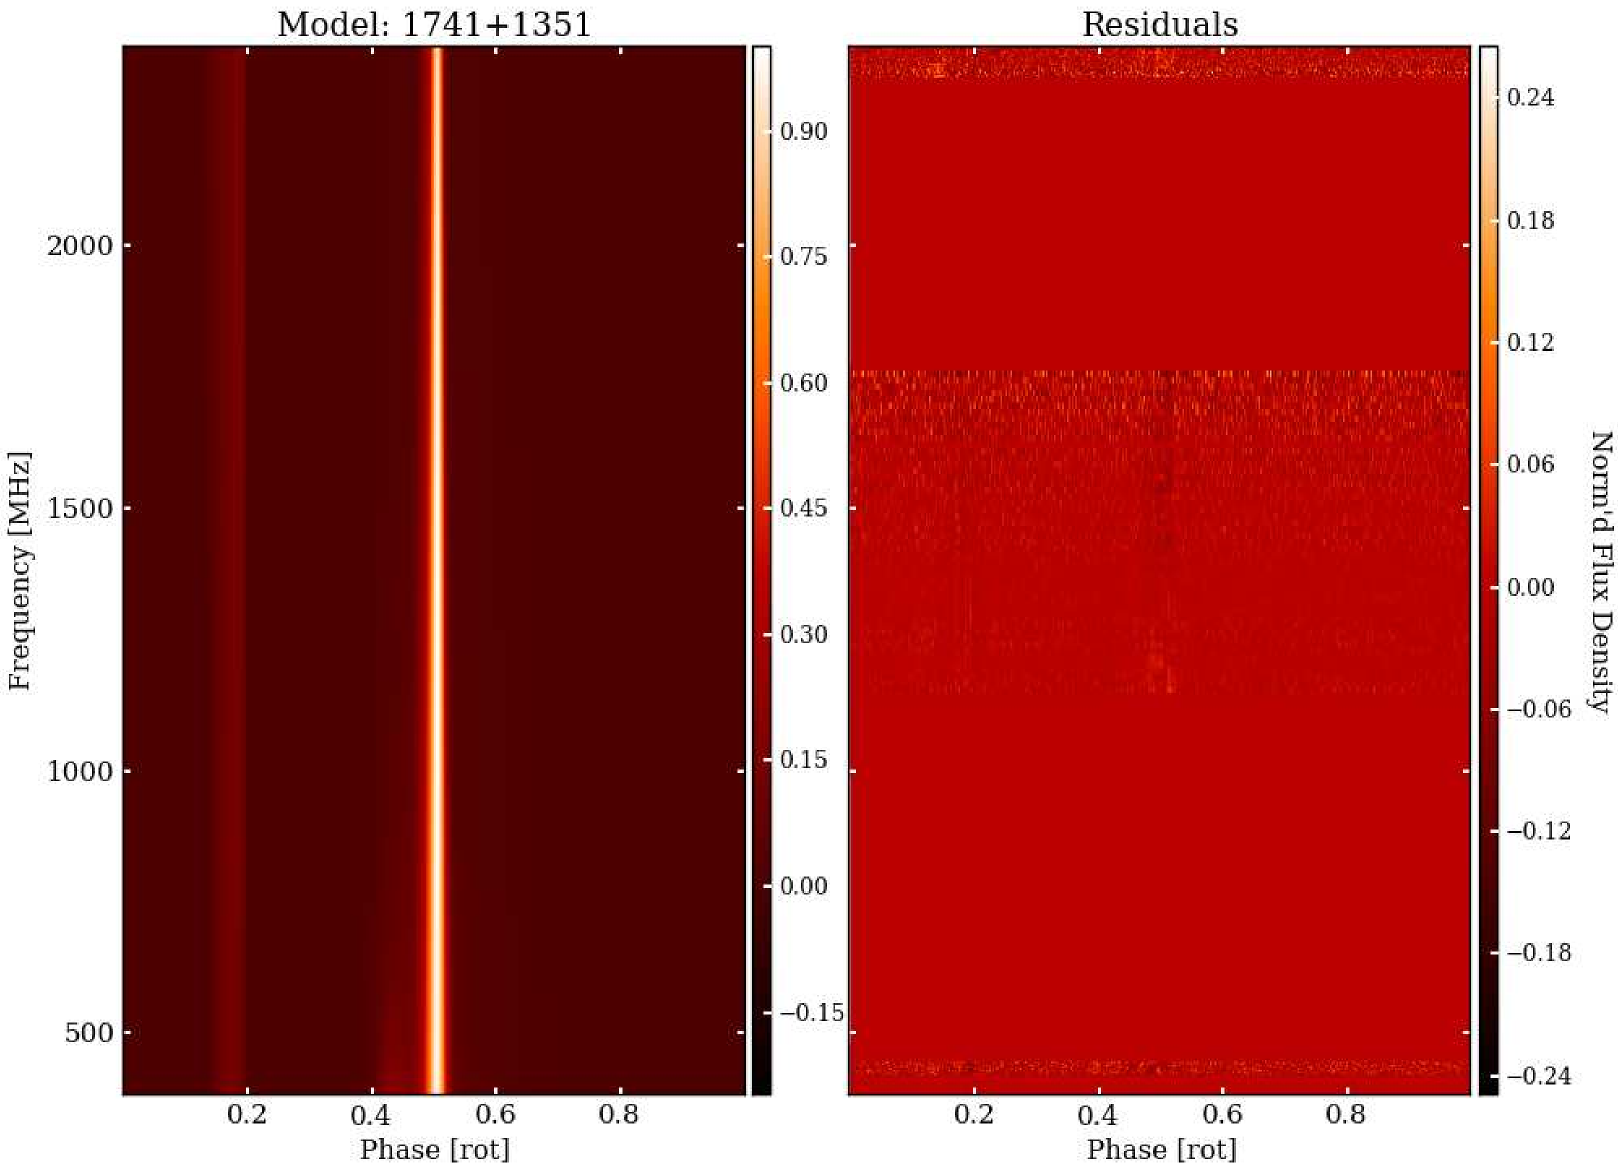

Fig. 4.187 – Fitted Gaussian component model and residuals for J1741+1351.

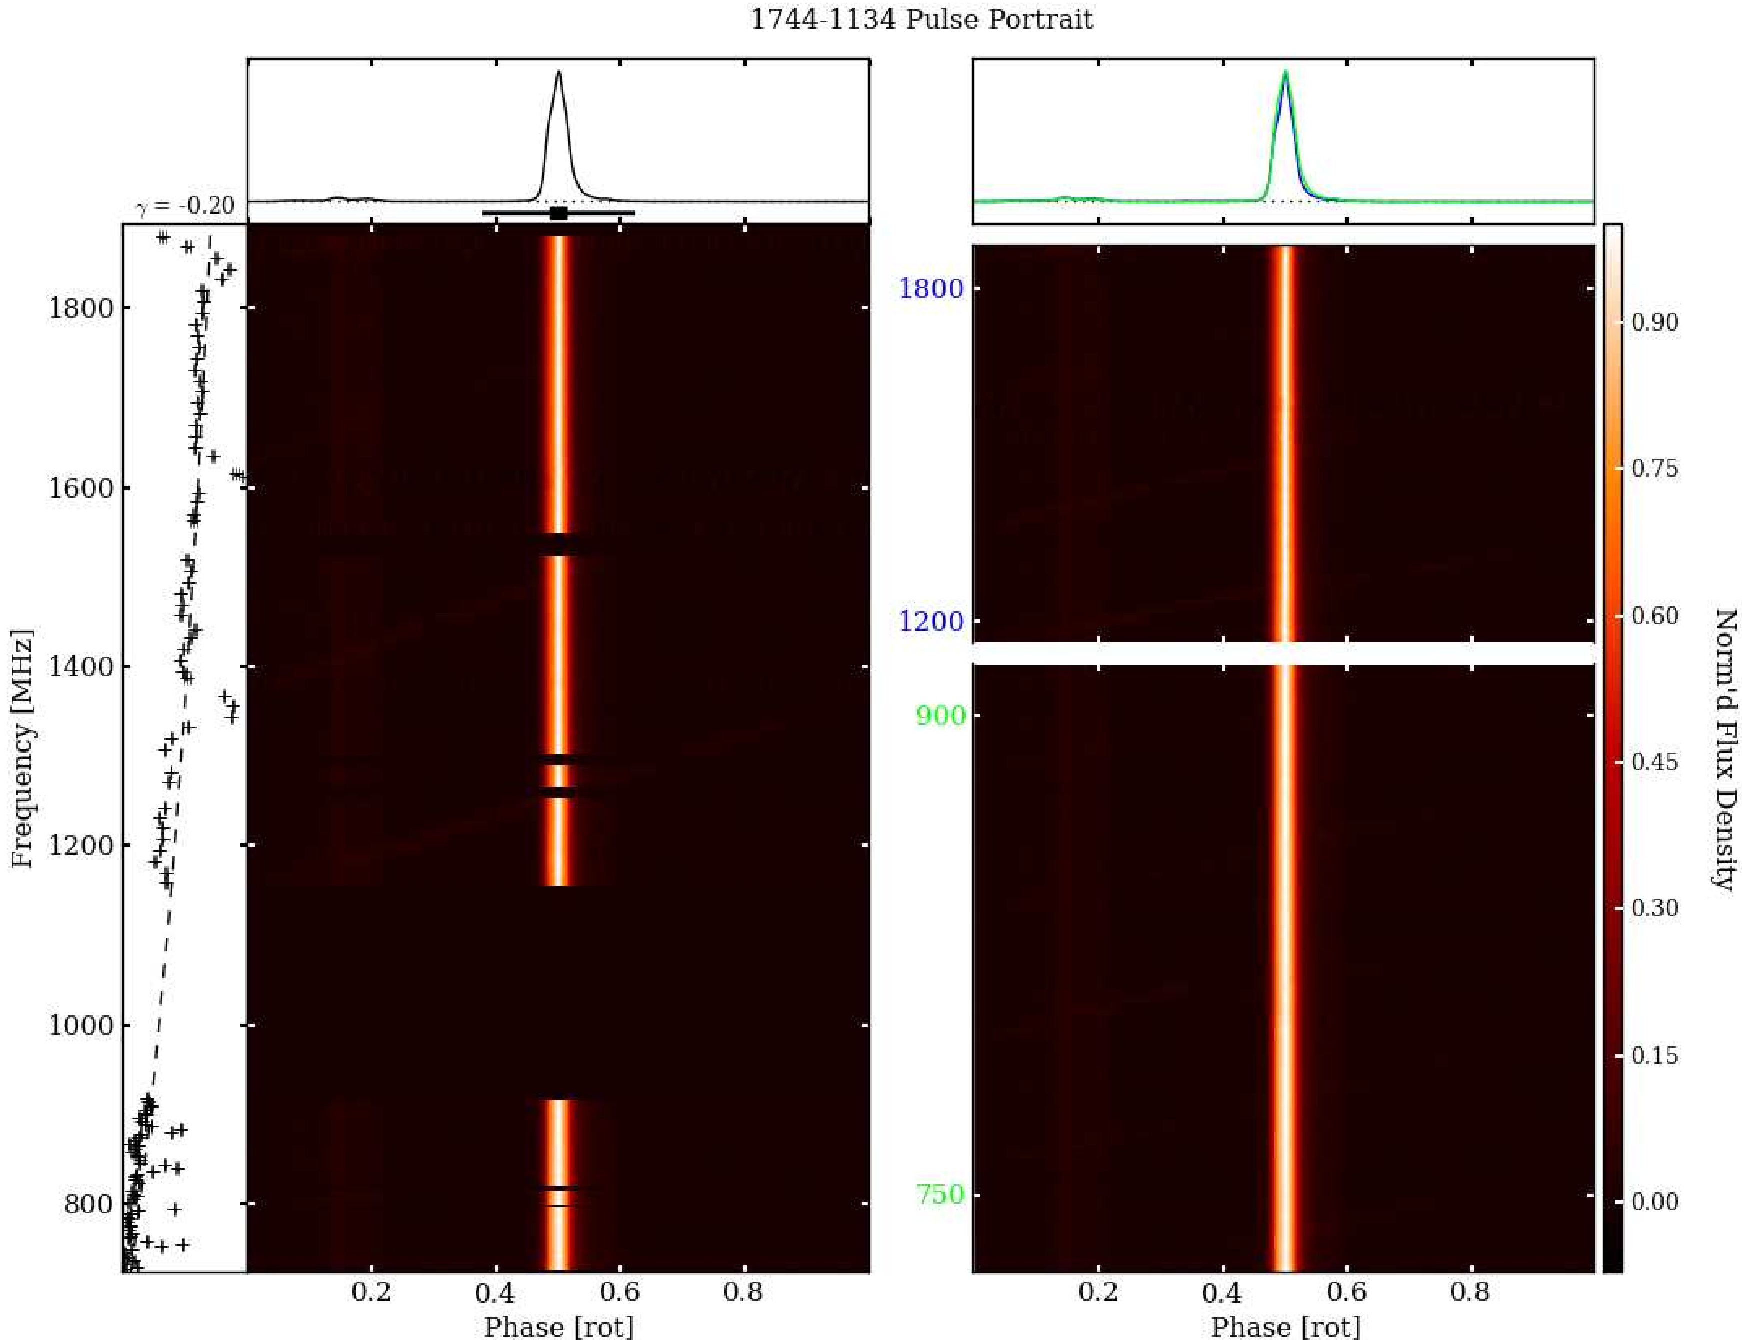

Fig. 4.188 – Concatenated portrait data and average profiles for J1744−1134.

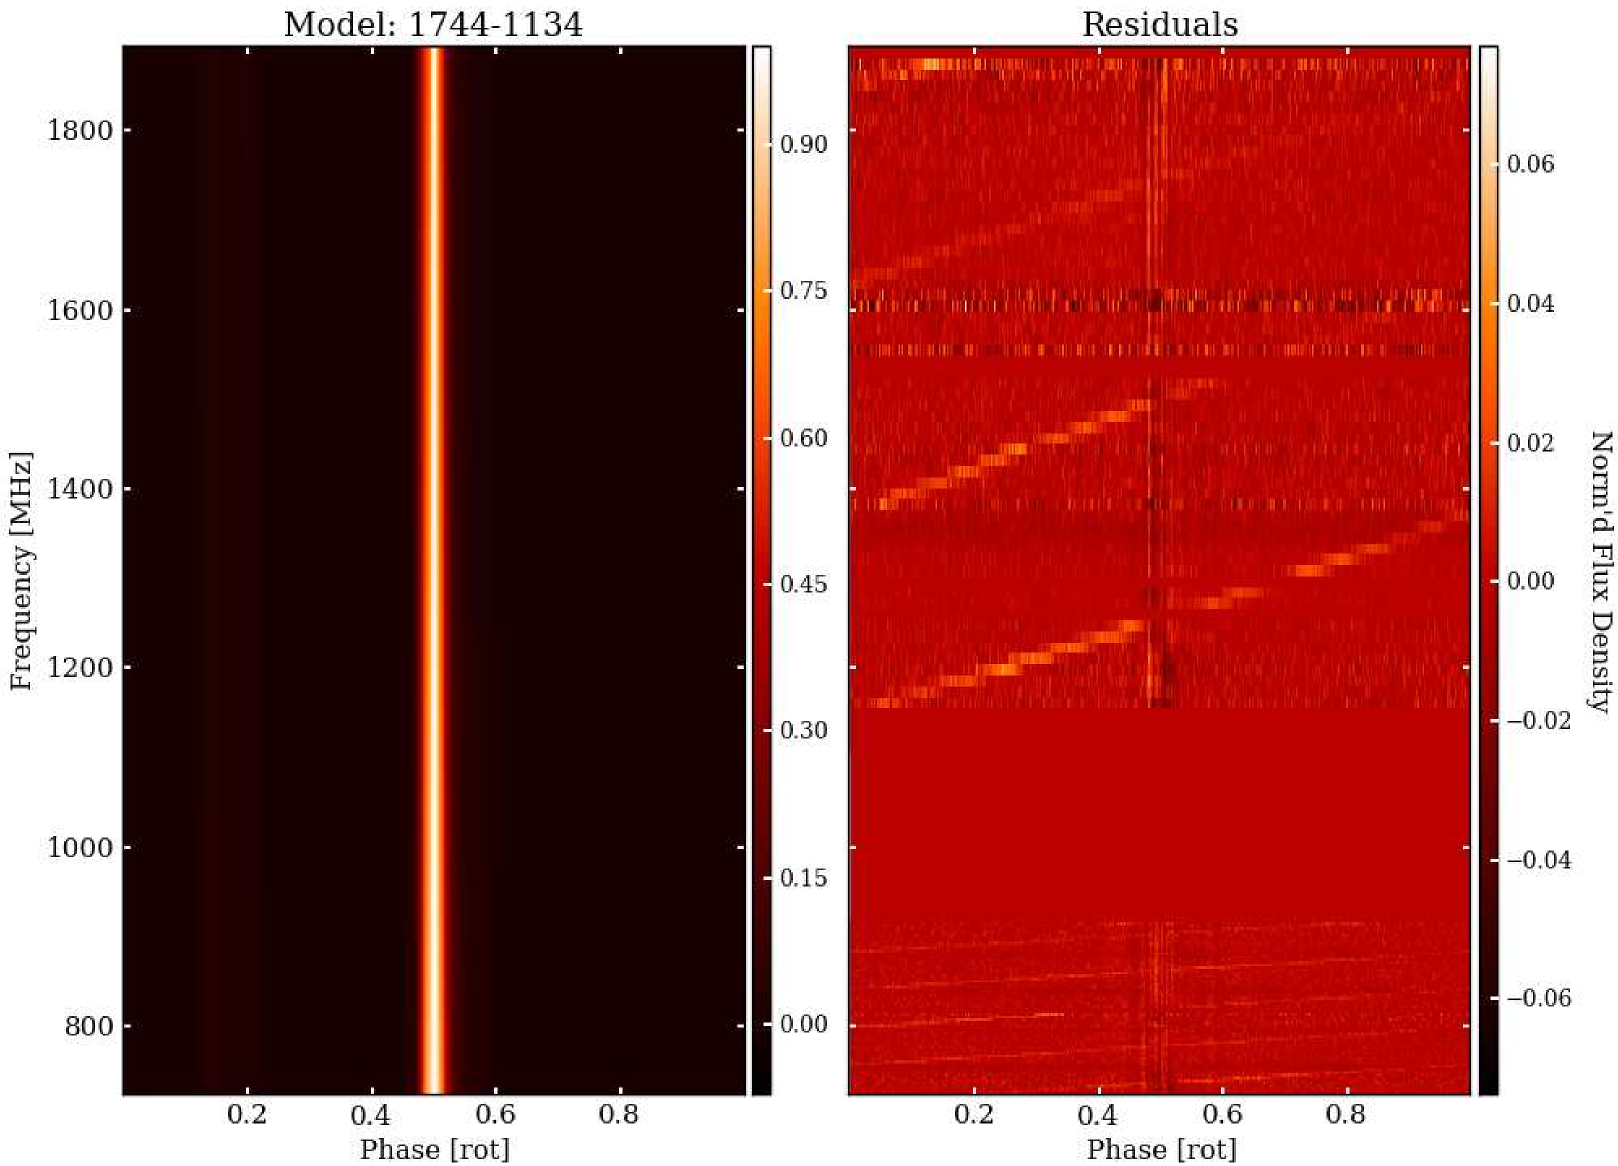

Fig. 4.189 – Fitted Gaussian component model and residuals for J1744−1134.

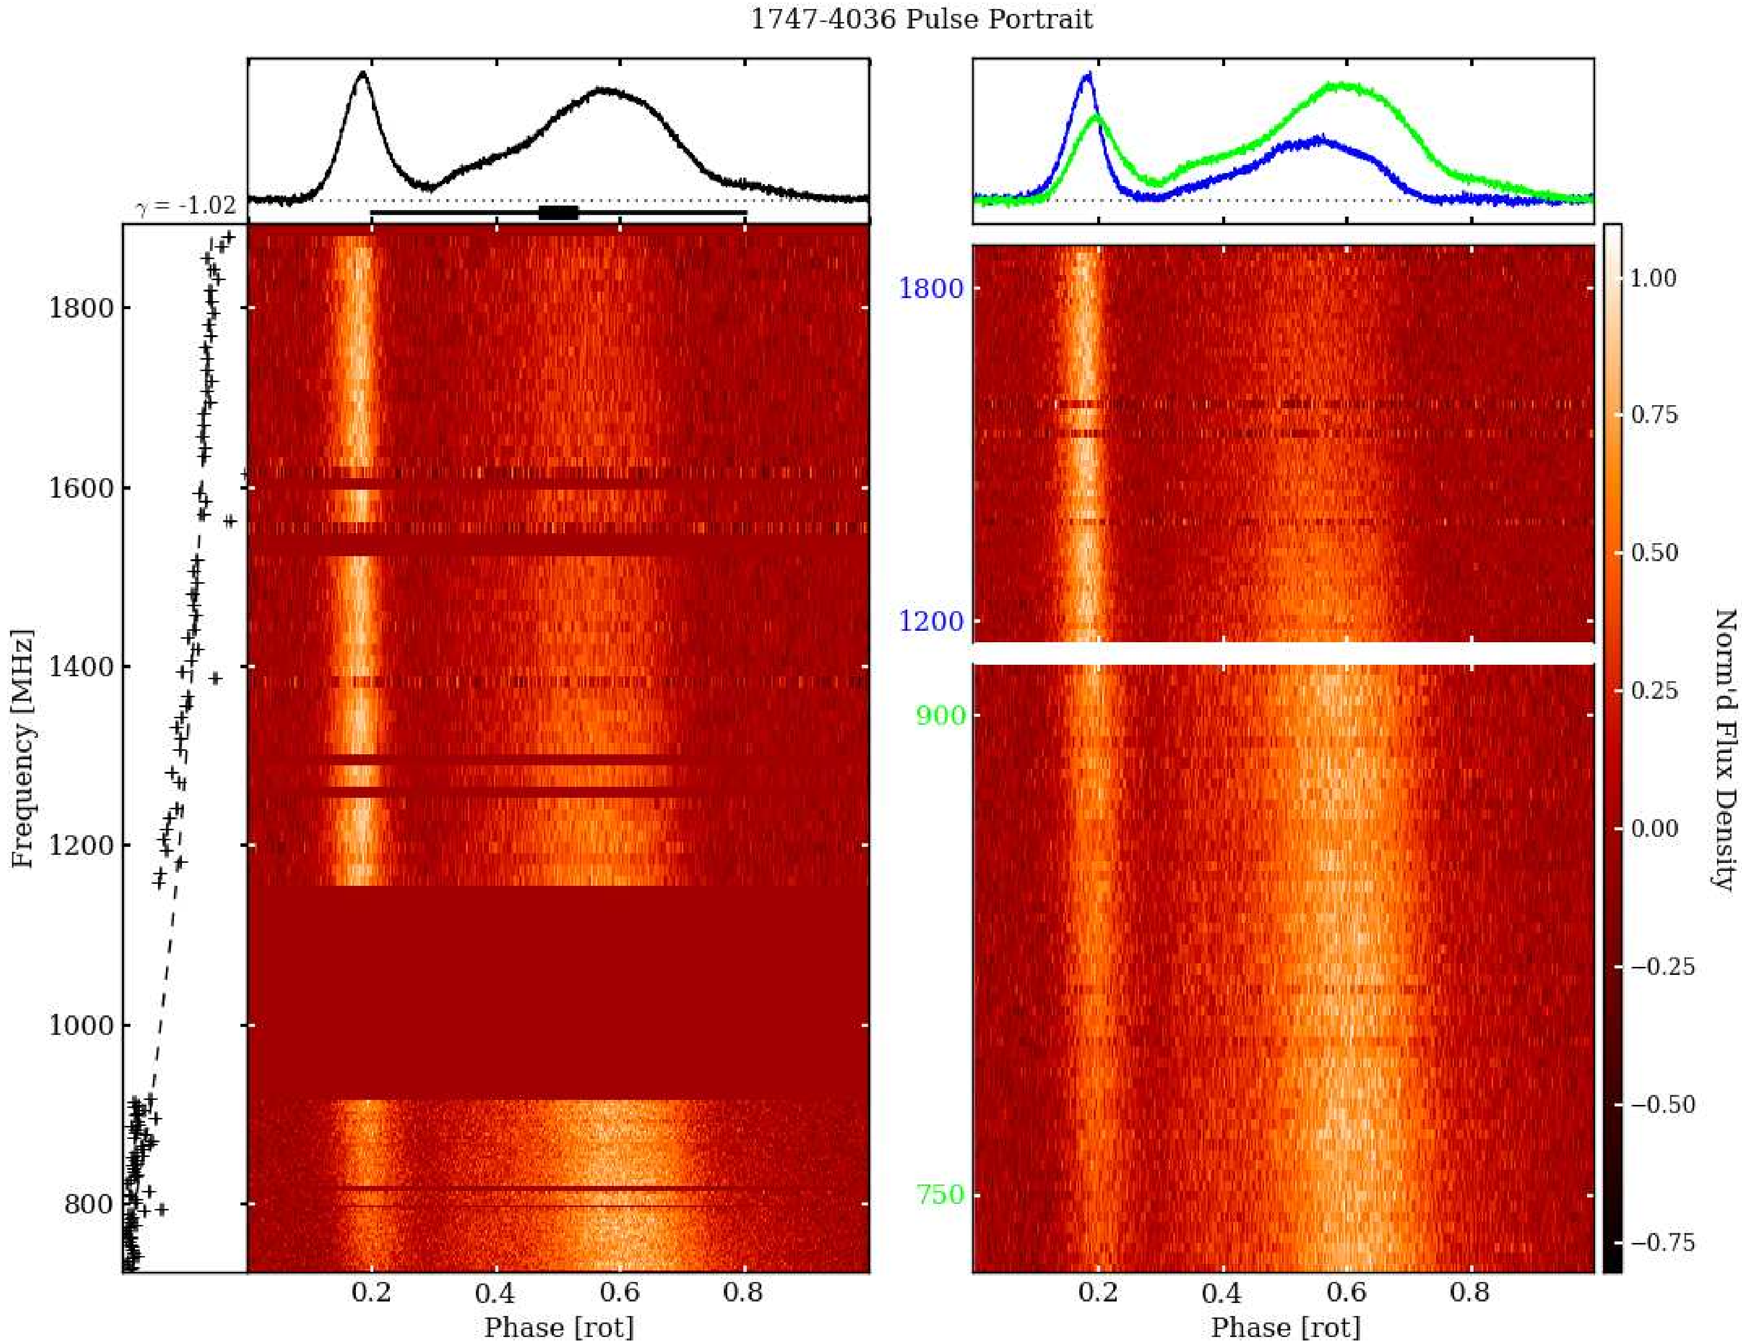

Fig. 4.190 – Concatenated portrait data and average profiles for J1747−4036.

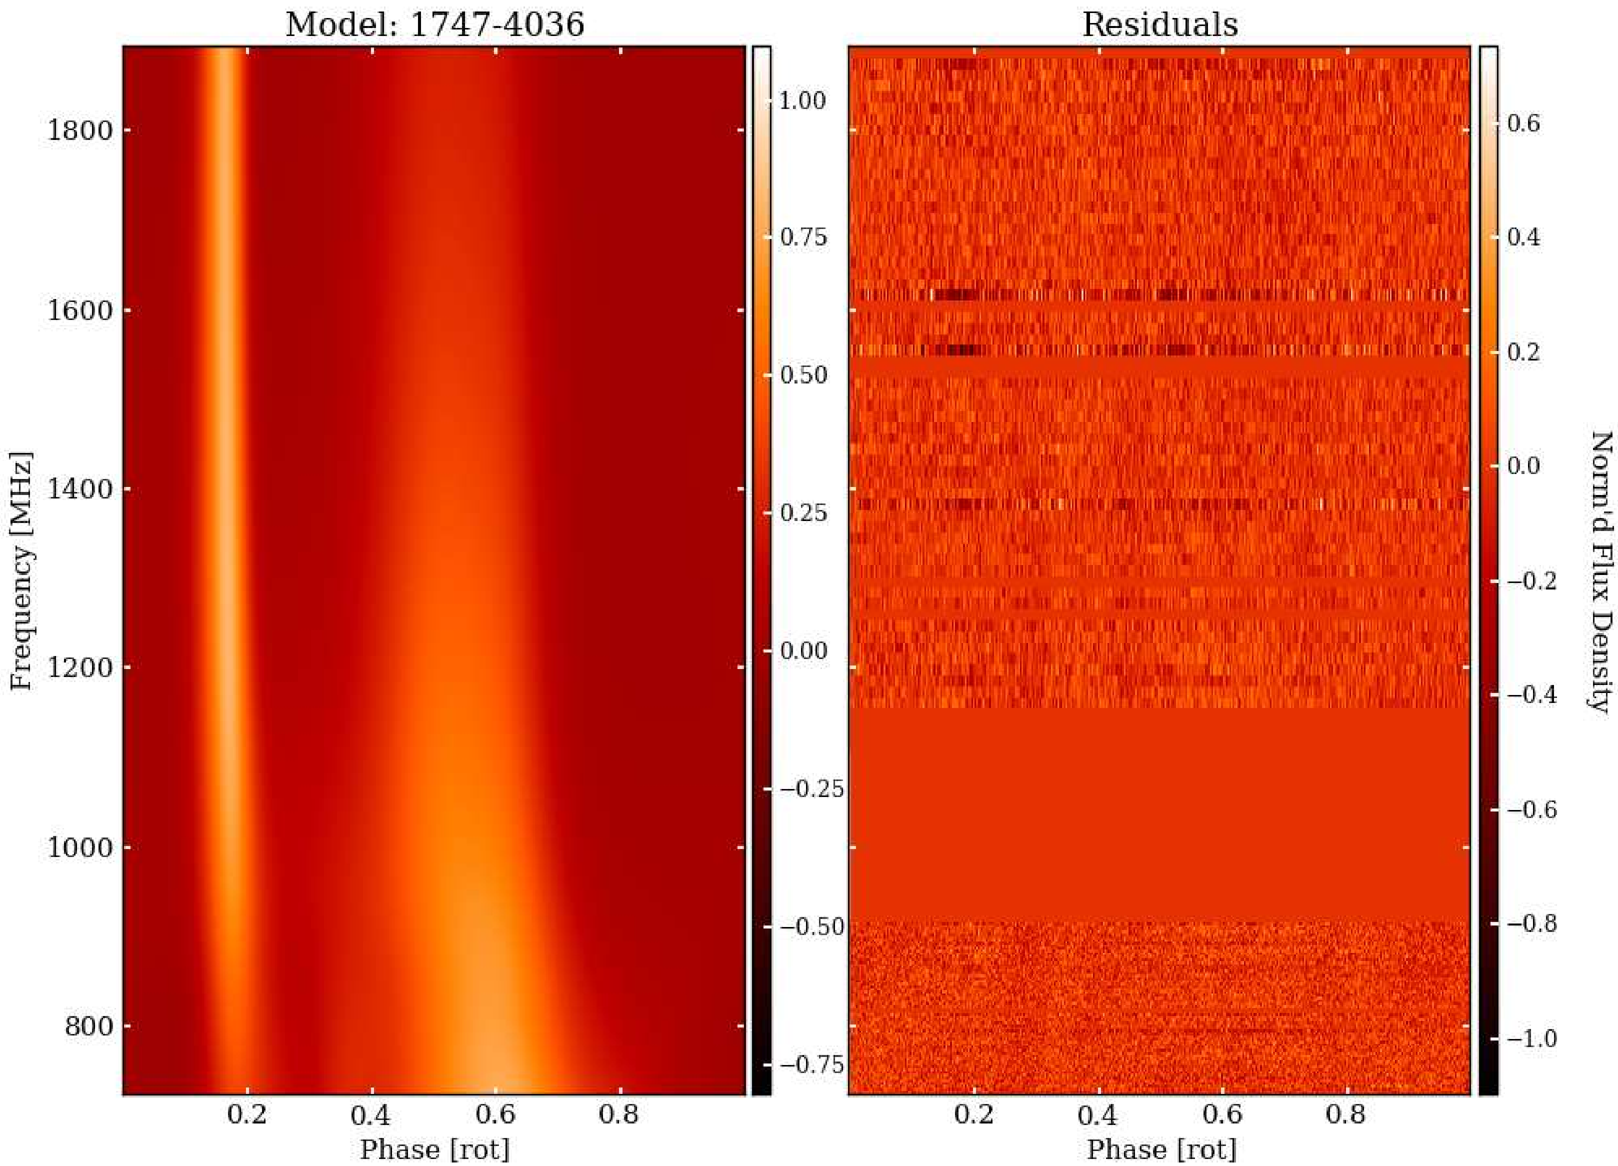

Fig. 4.191 – Fitted Gaussian component model and residuals for J1747−4036.

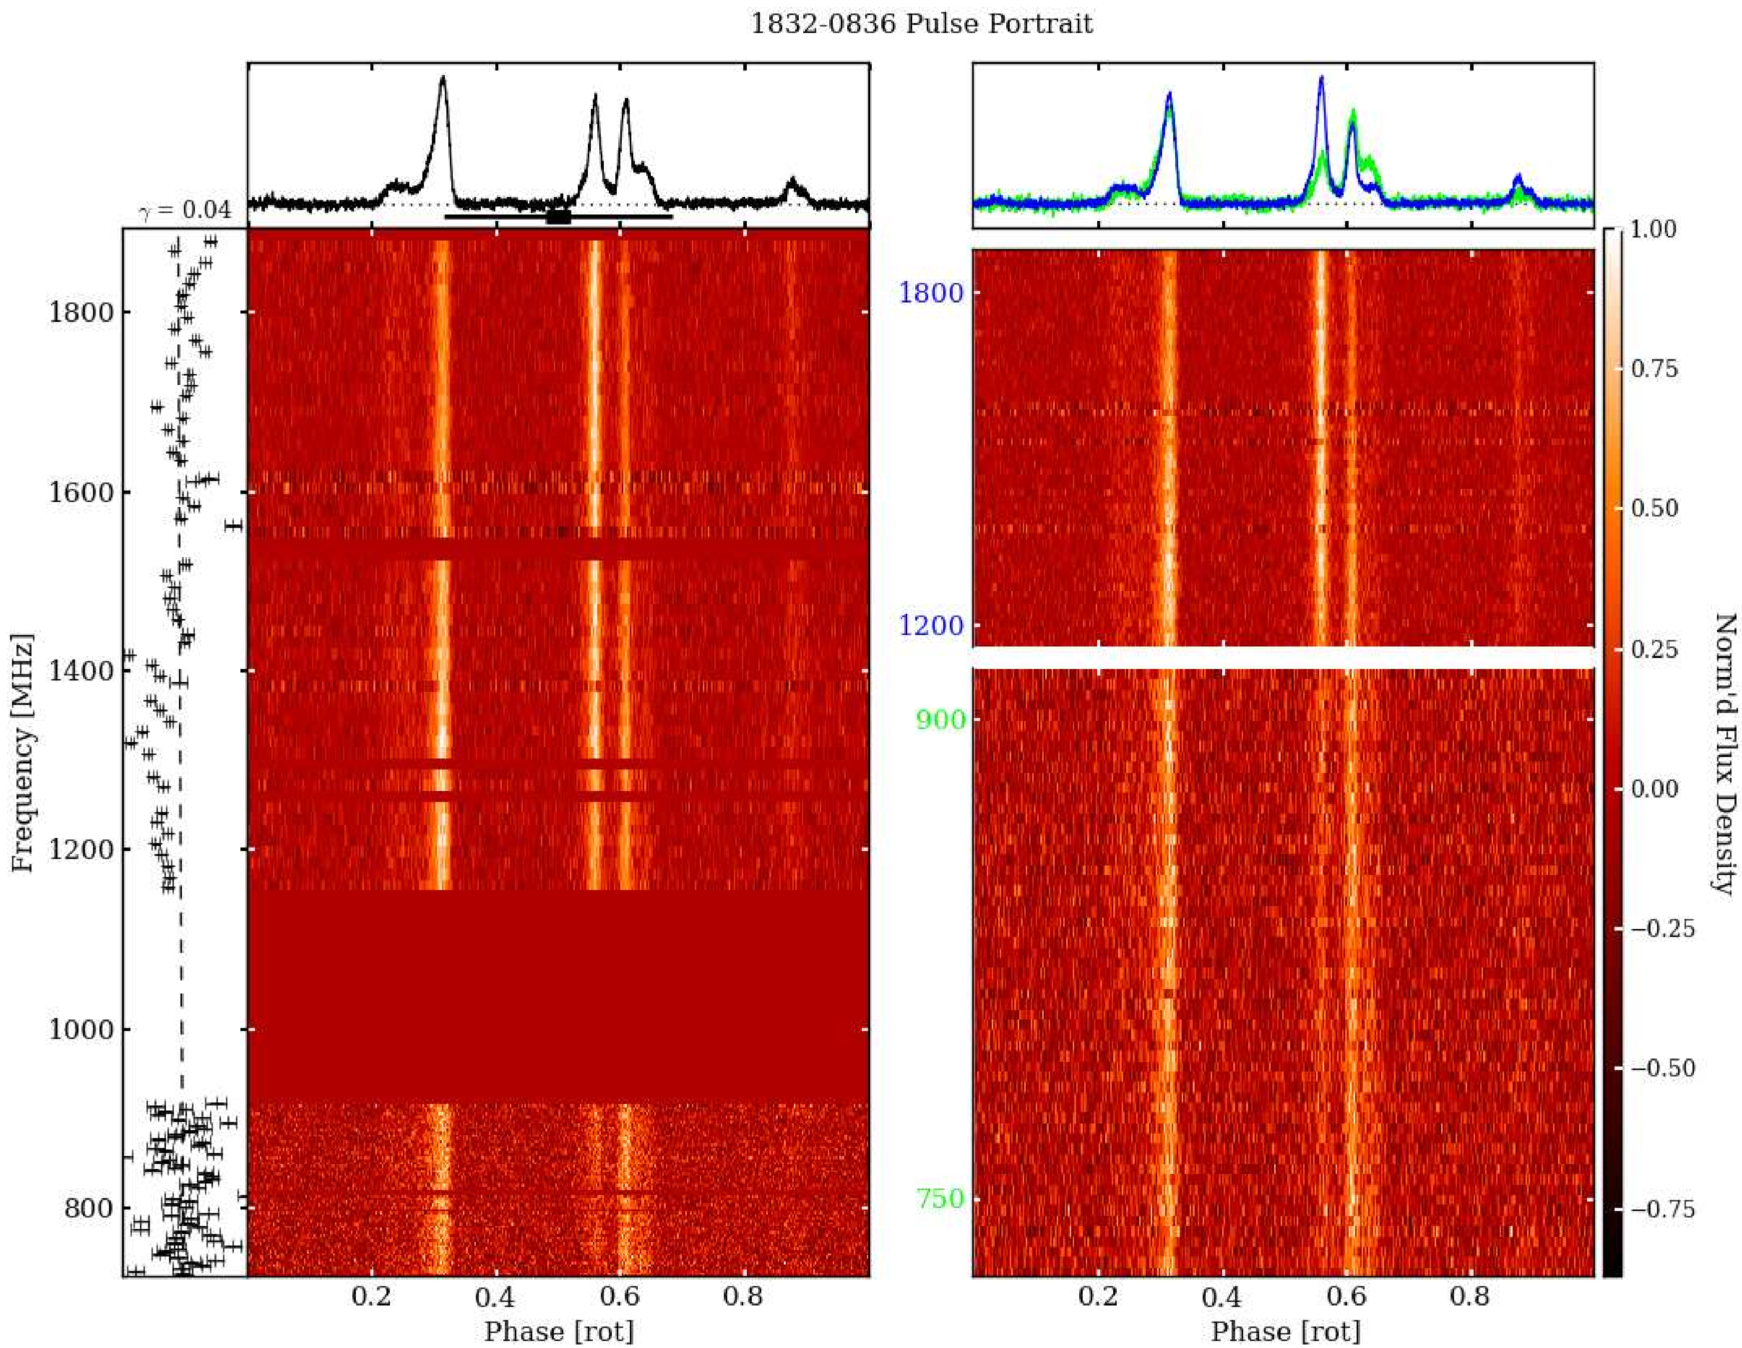

Fig. 4.192 – Concatenated portrait data and average profiles for J1832−0836.

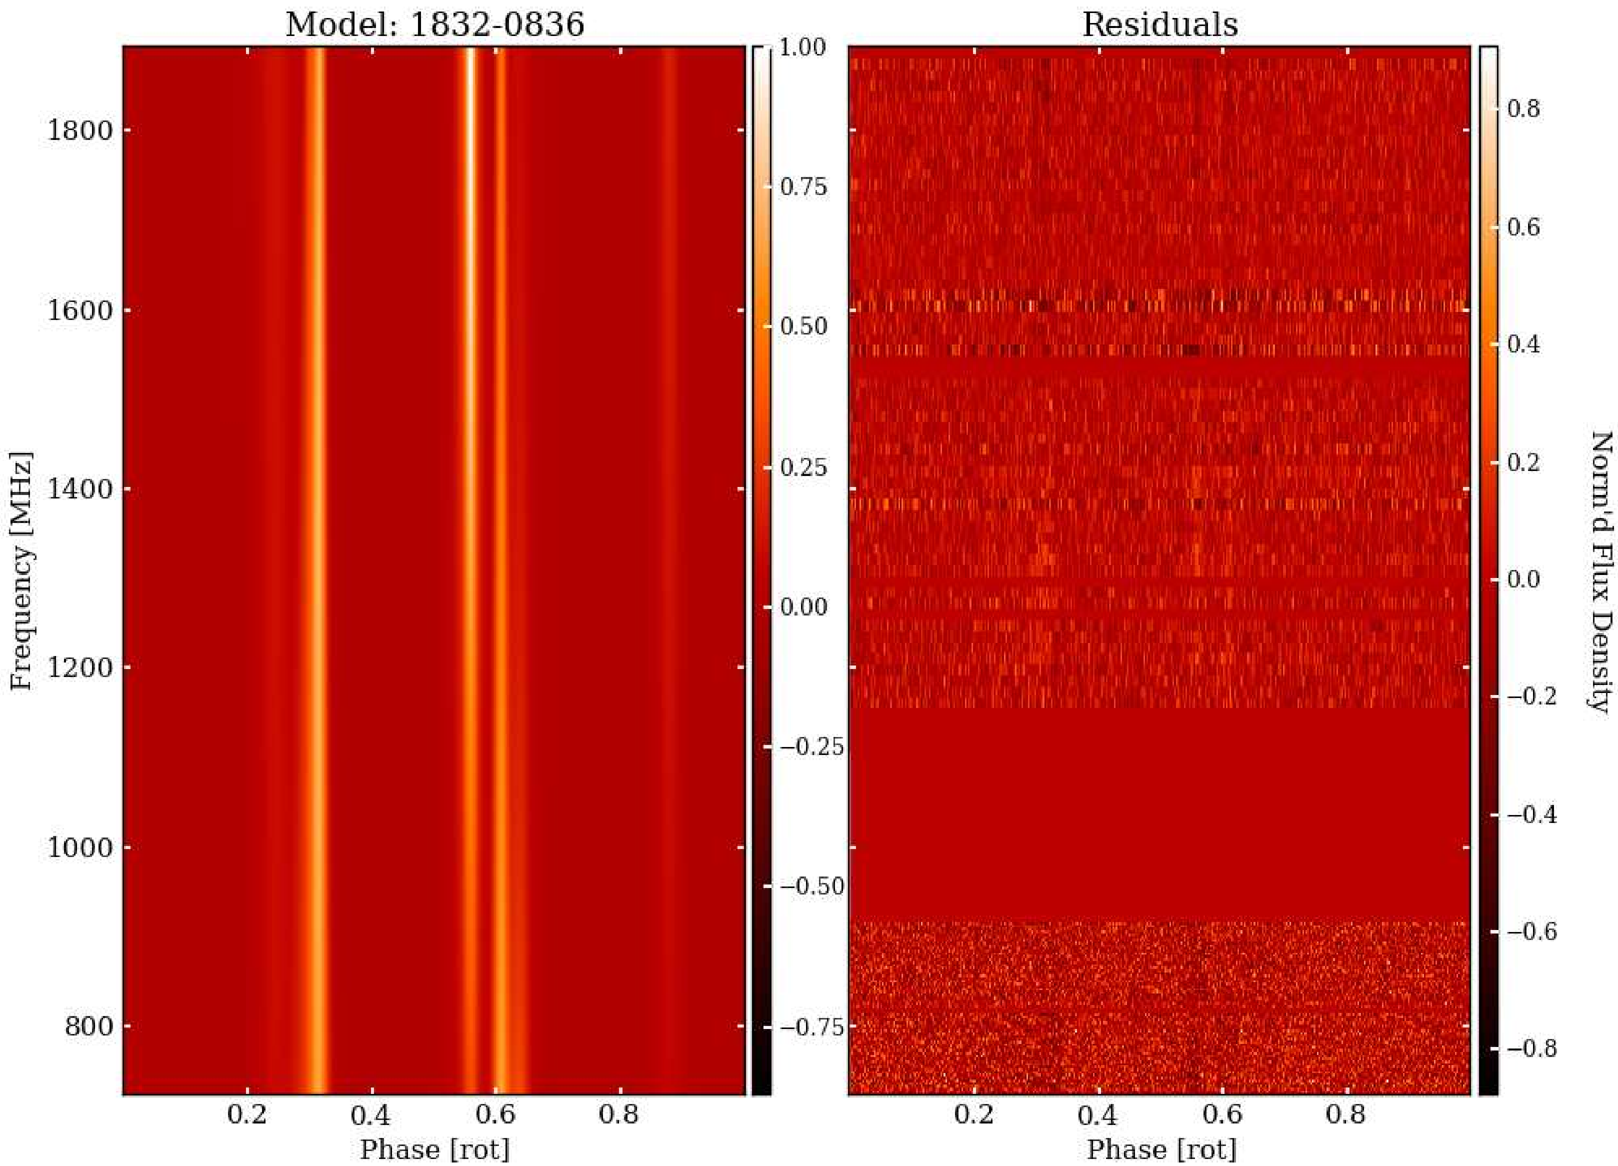

Fig. 4.193 – Fitted Gaussian component model and residuals for J1832−0836.

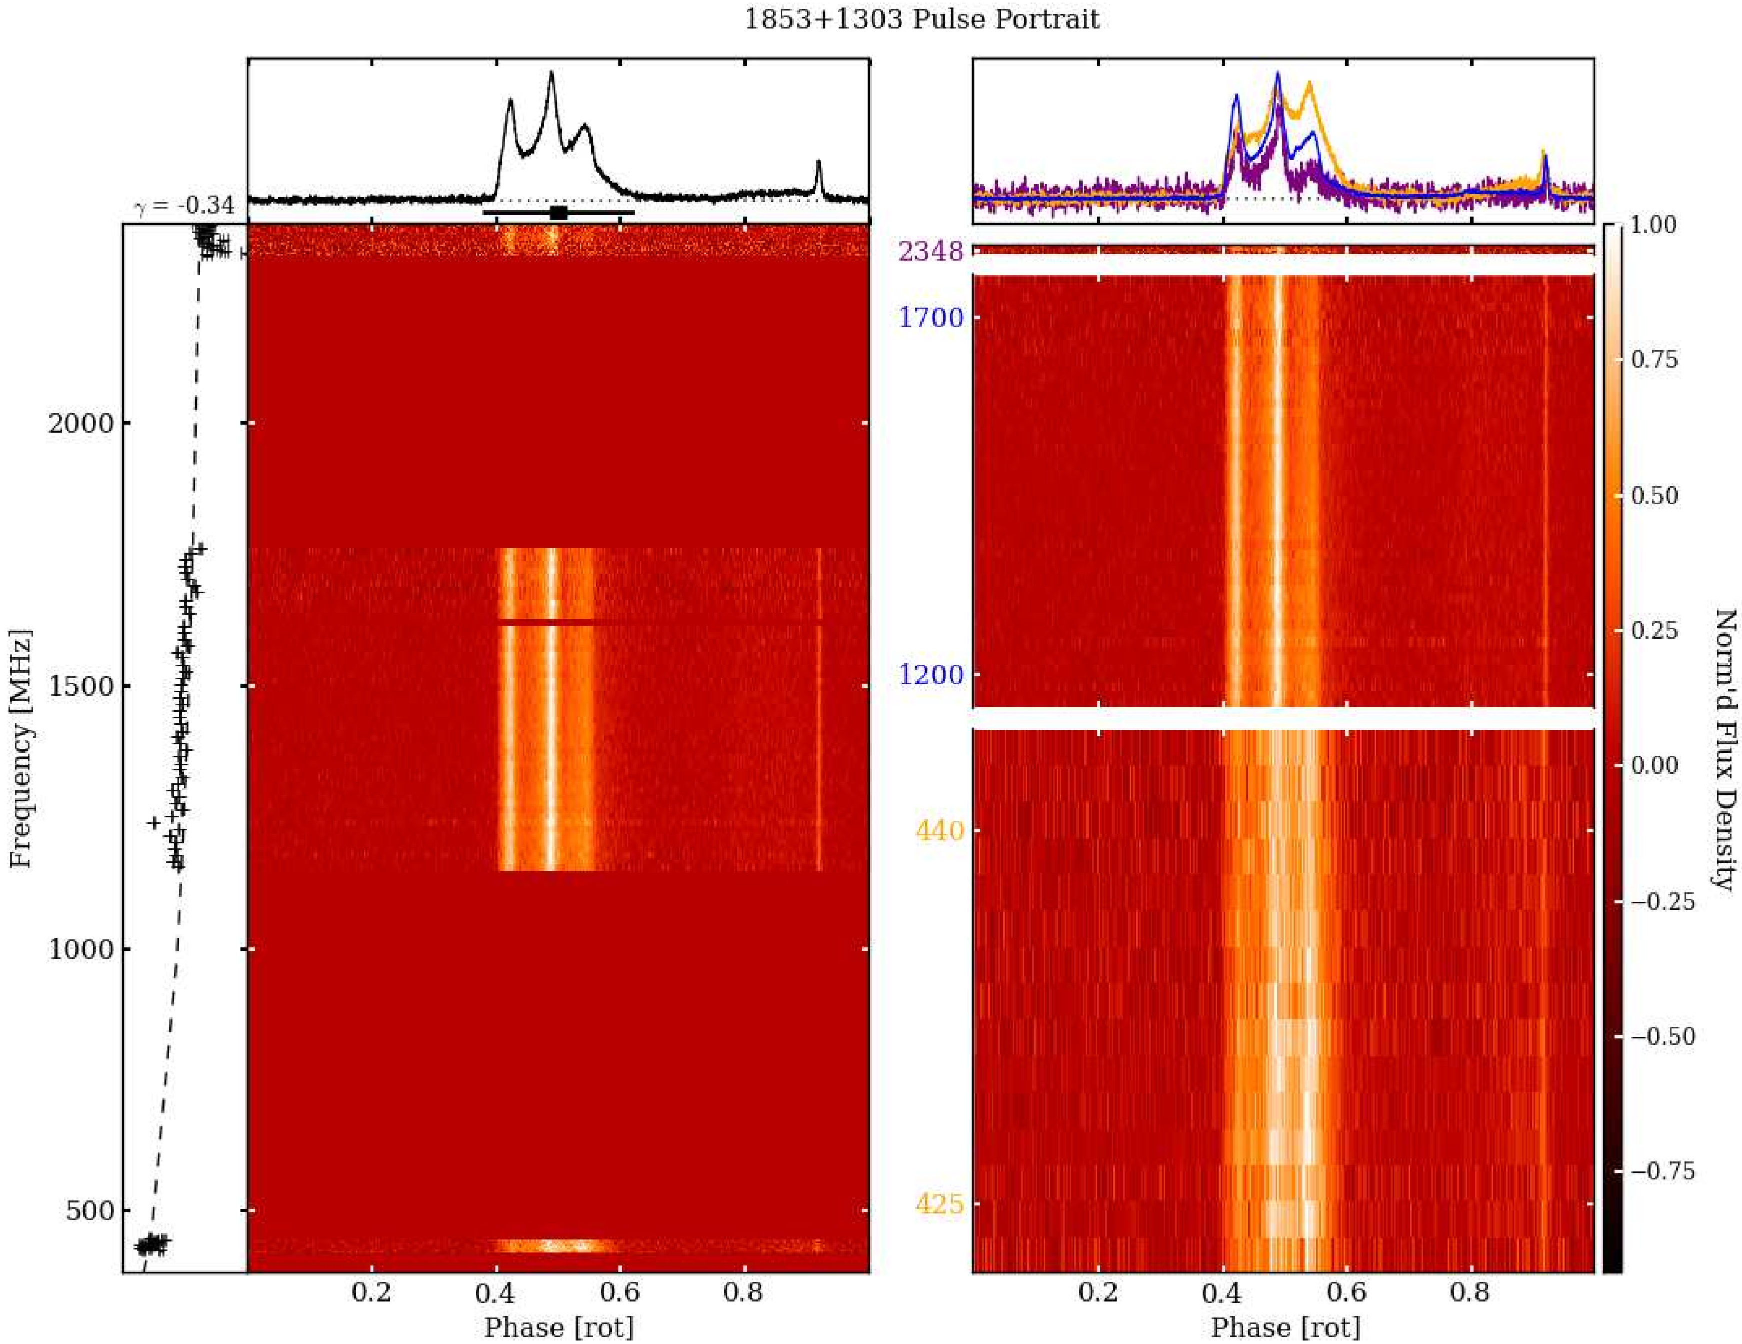

Fig. 4.194 – Concatenated portrait data and average profiles for J1853+1303.

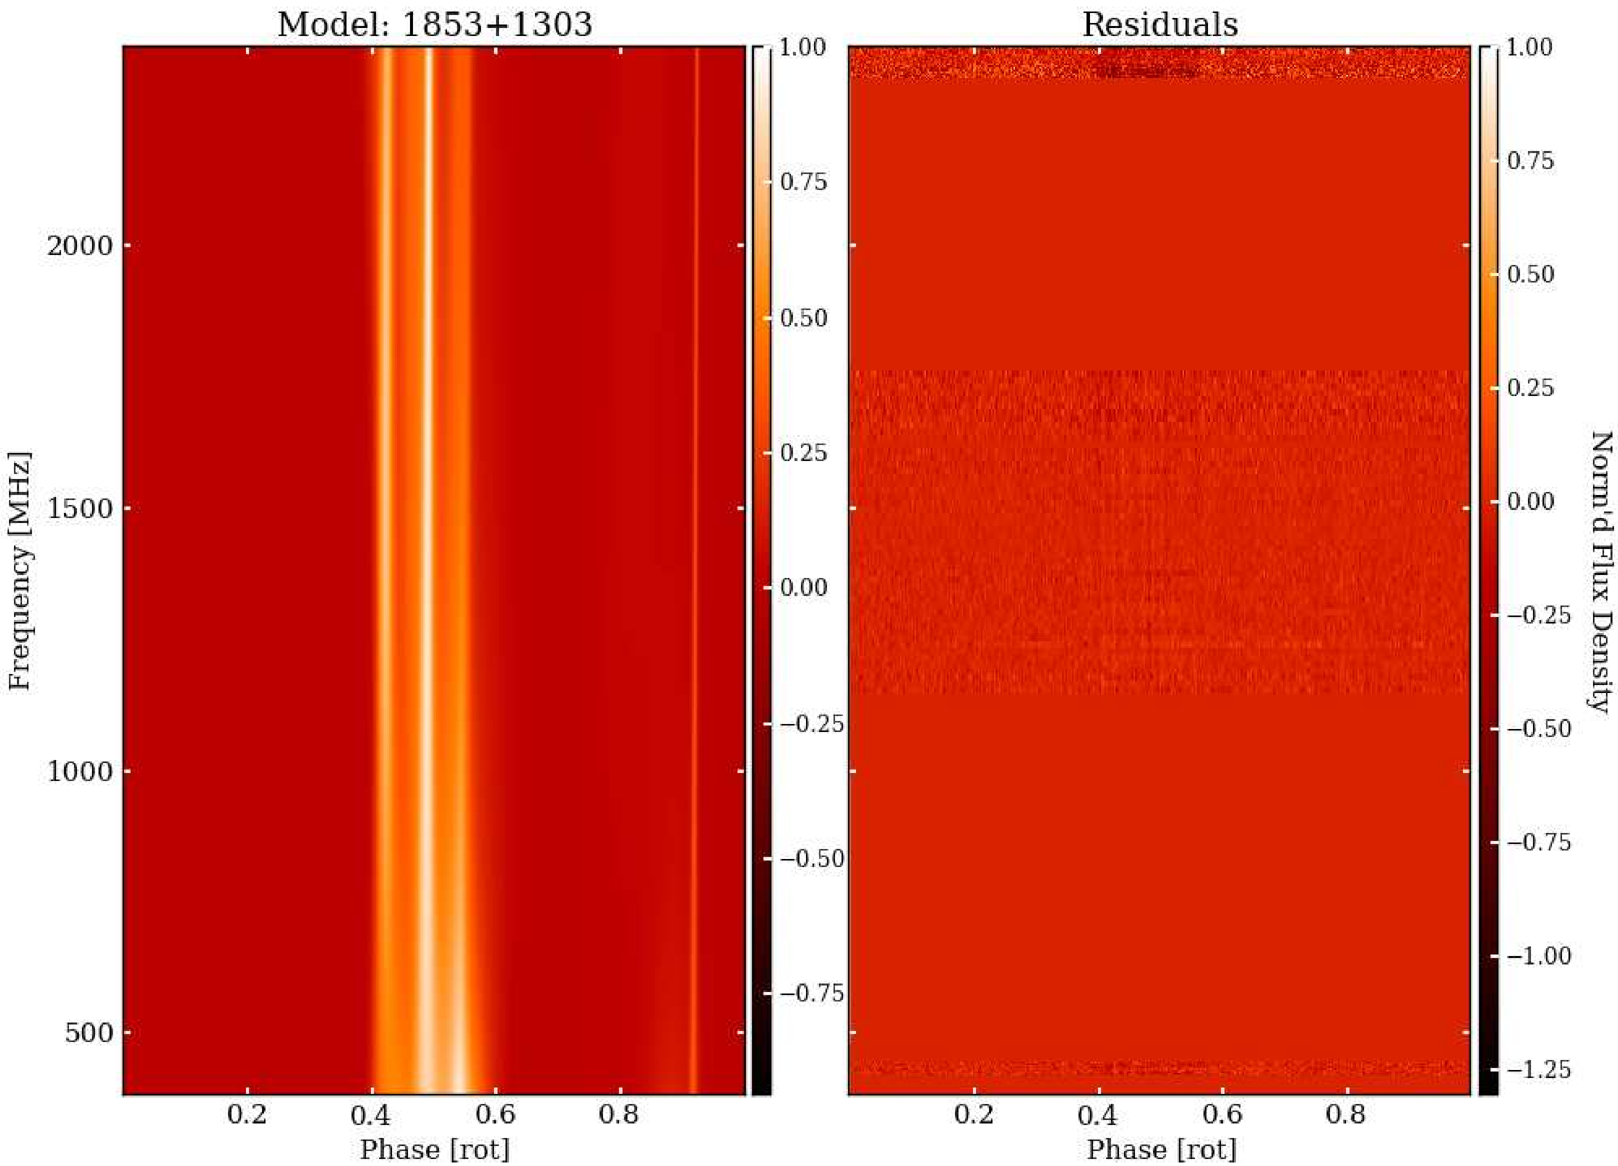

Fig. 4.195 – Fitted Gaussian component model and residuals for J1853+1303.

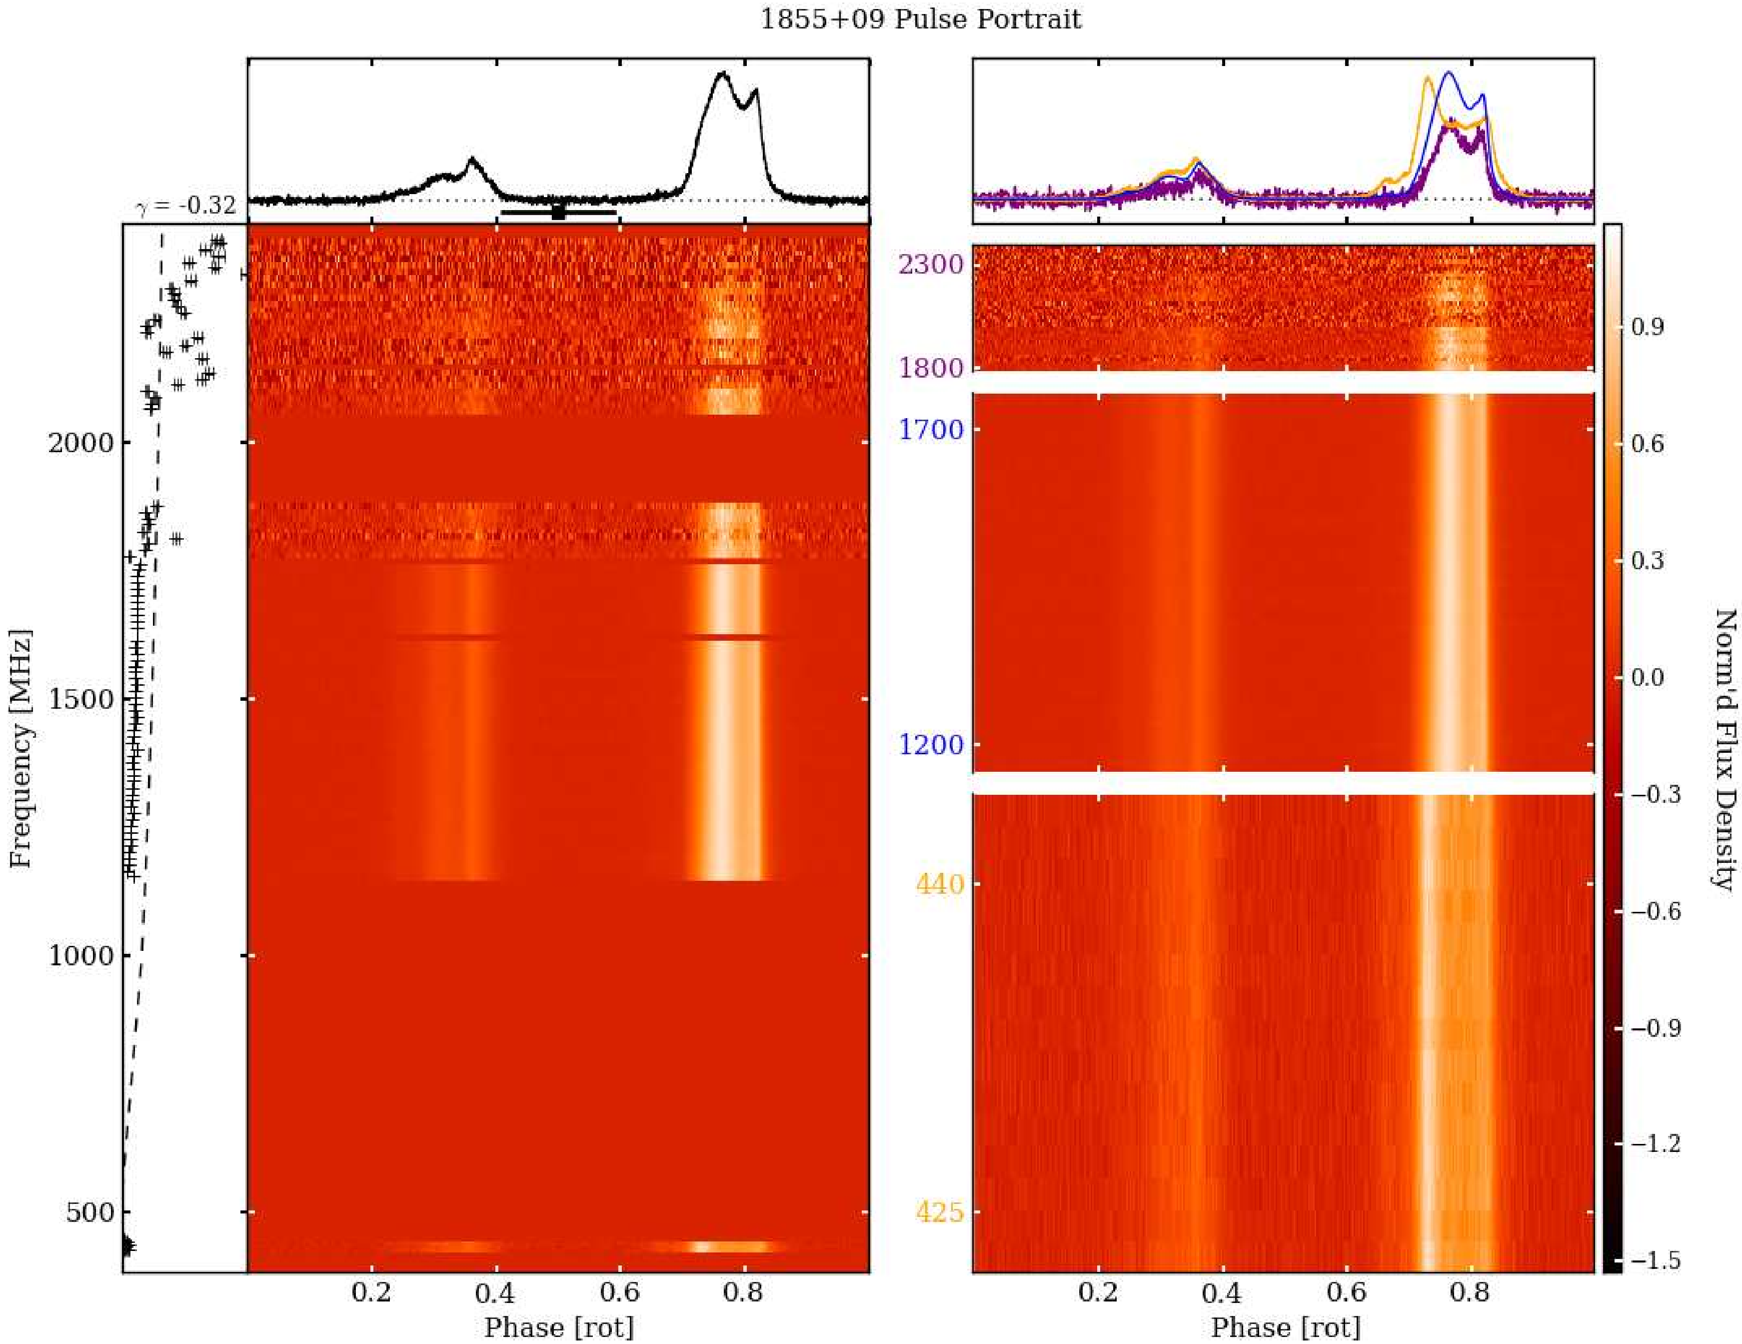

Fig. 4.196 – Concatenated portrait data and average profiles for B1855+09.

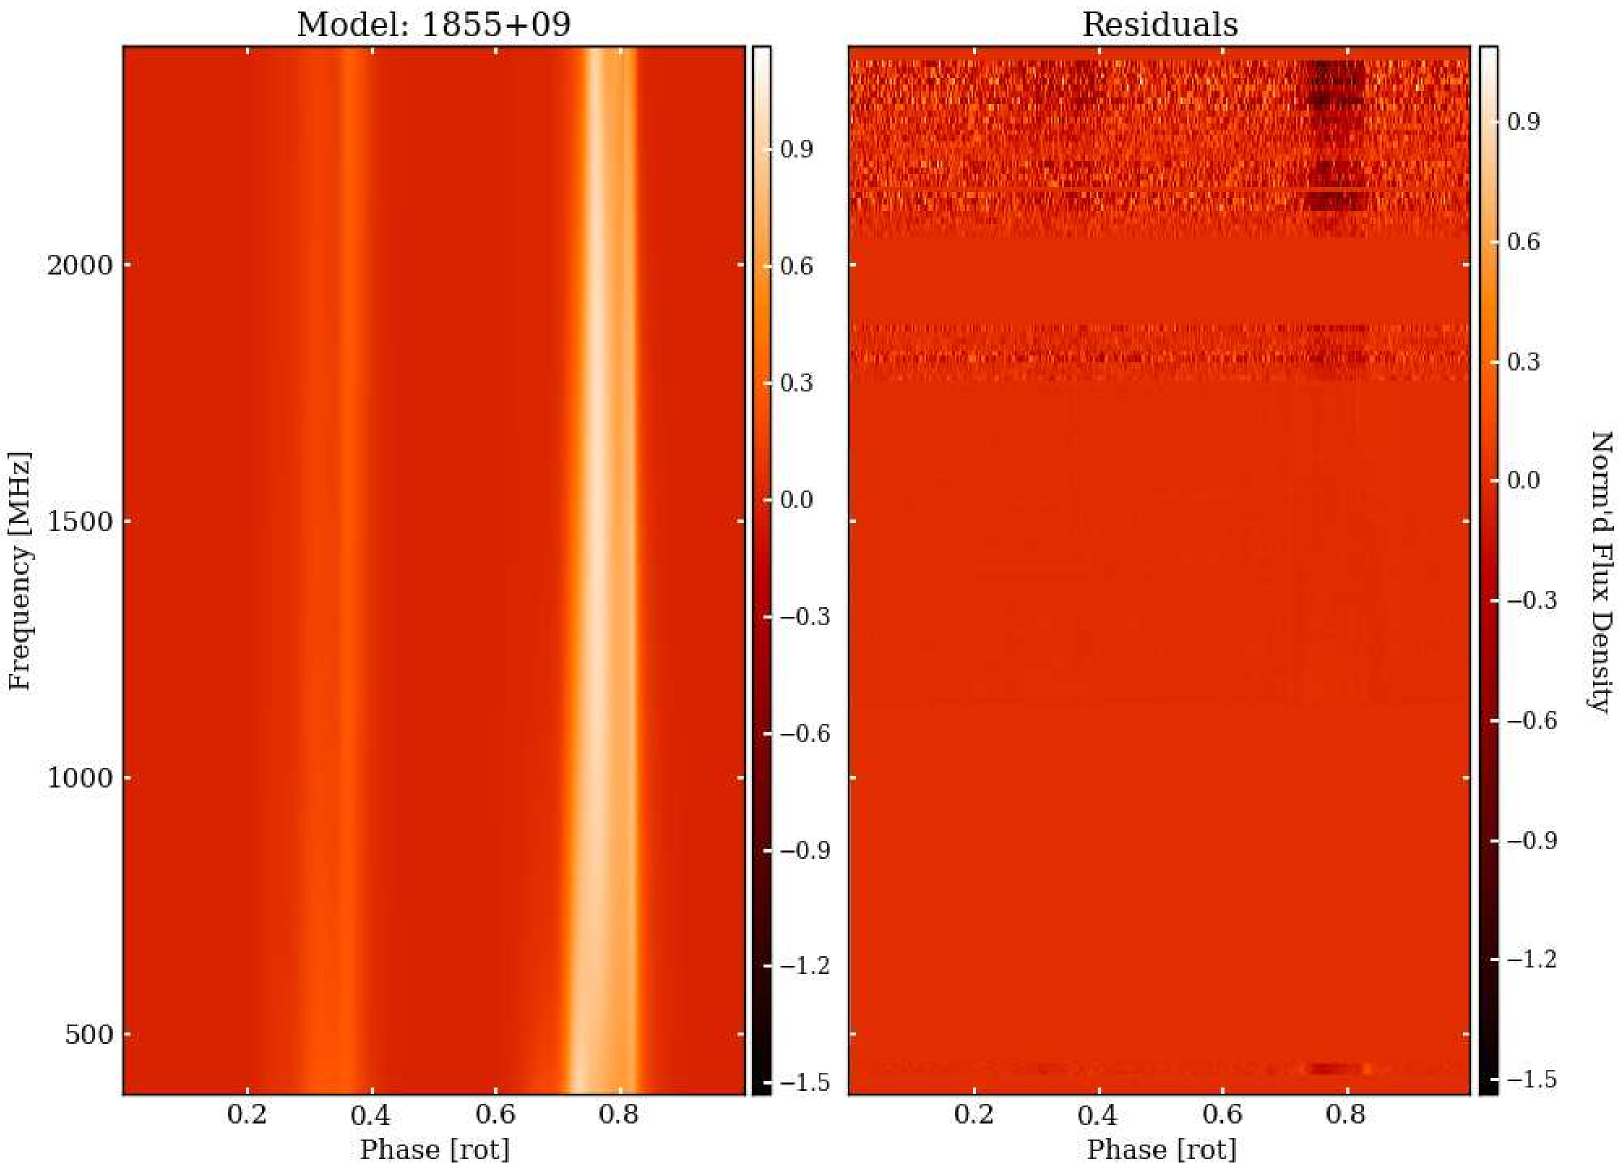

Fig. 4.197 – Fitted Gaussian component model and residuals for B1855+09.

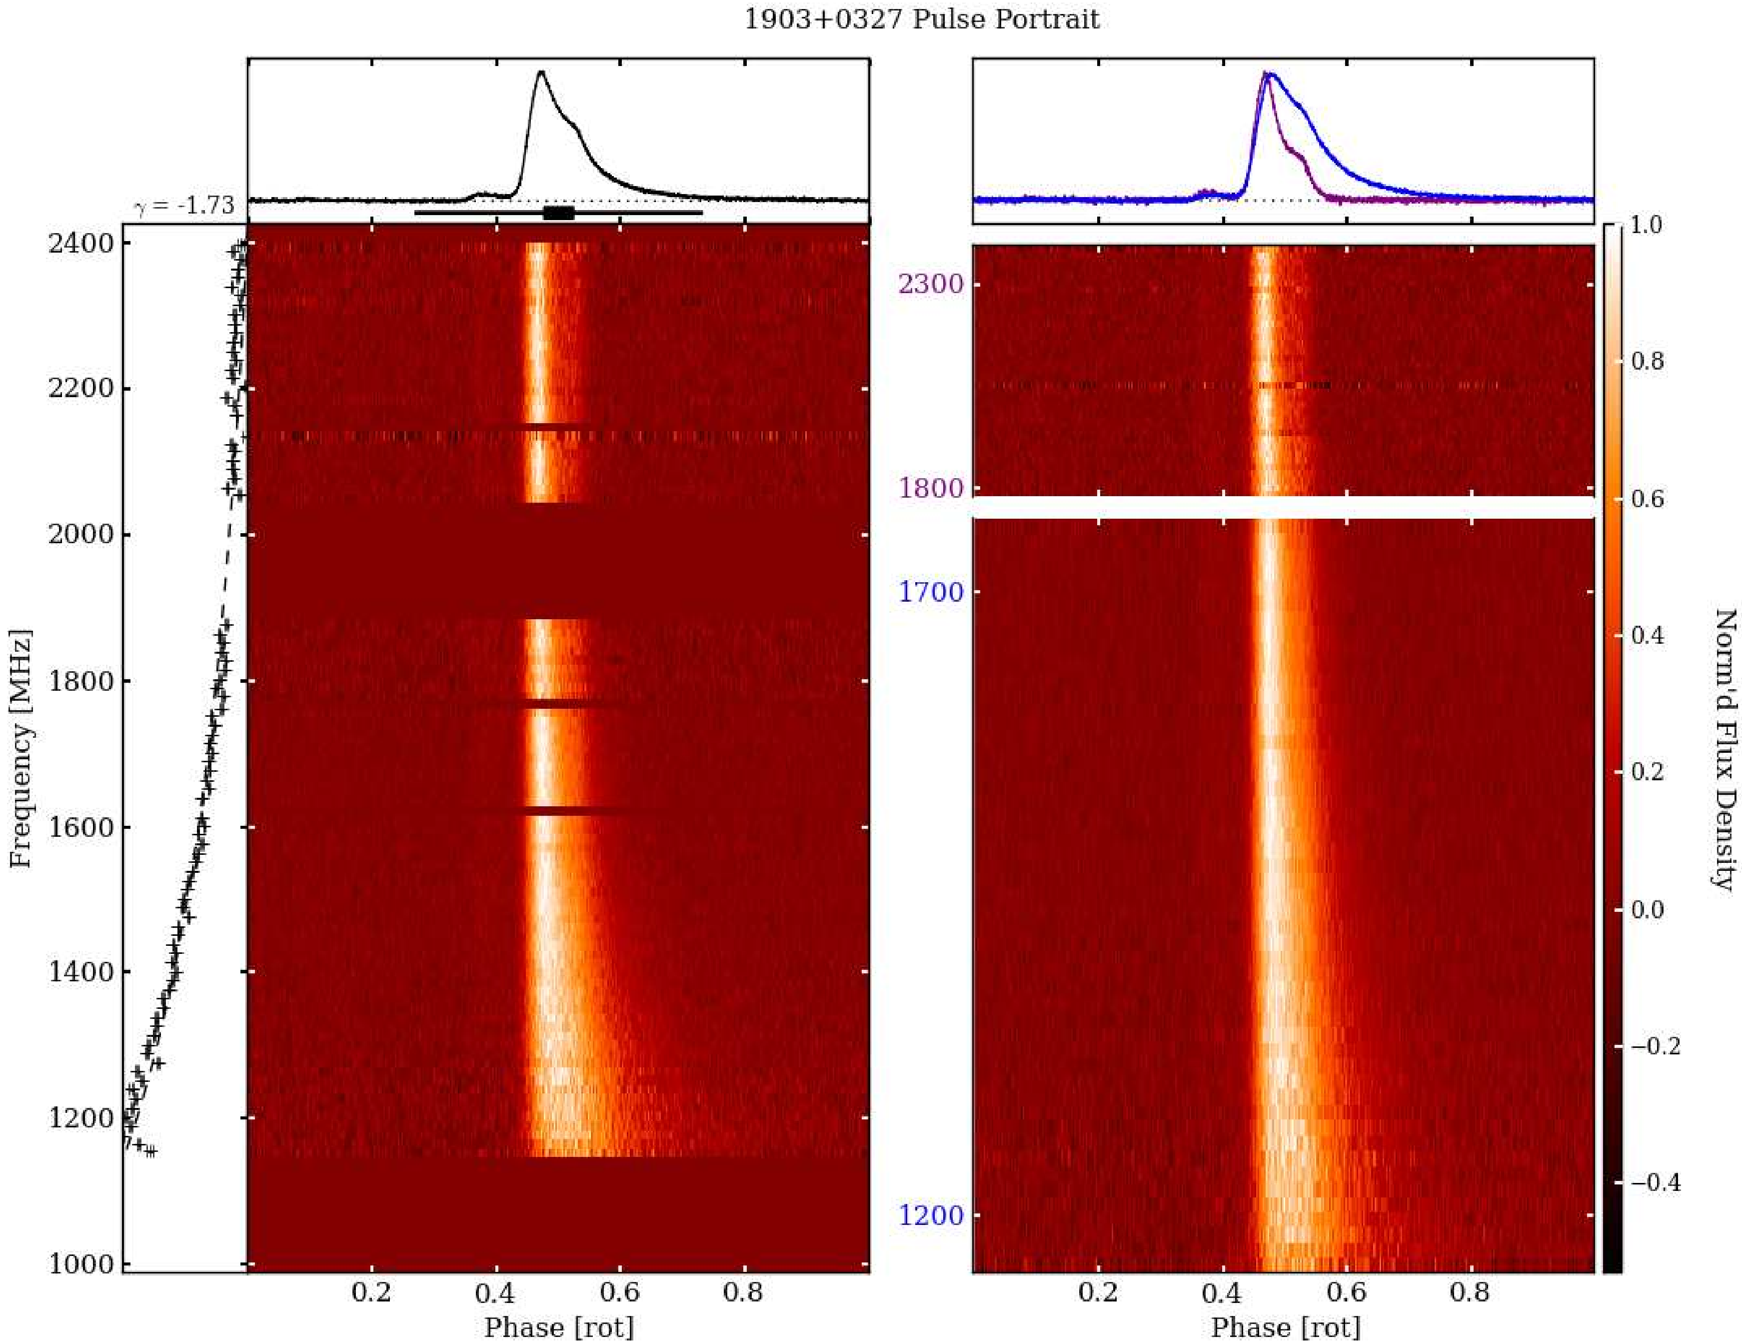

Fig. 4.198 – Concatenated portrait data and average profiles for J1903+0327.

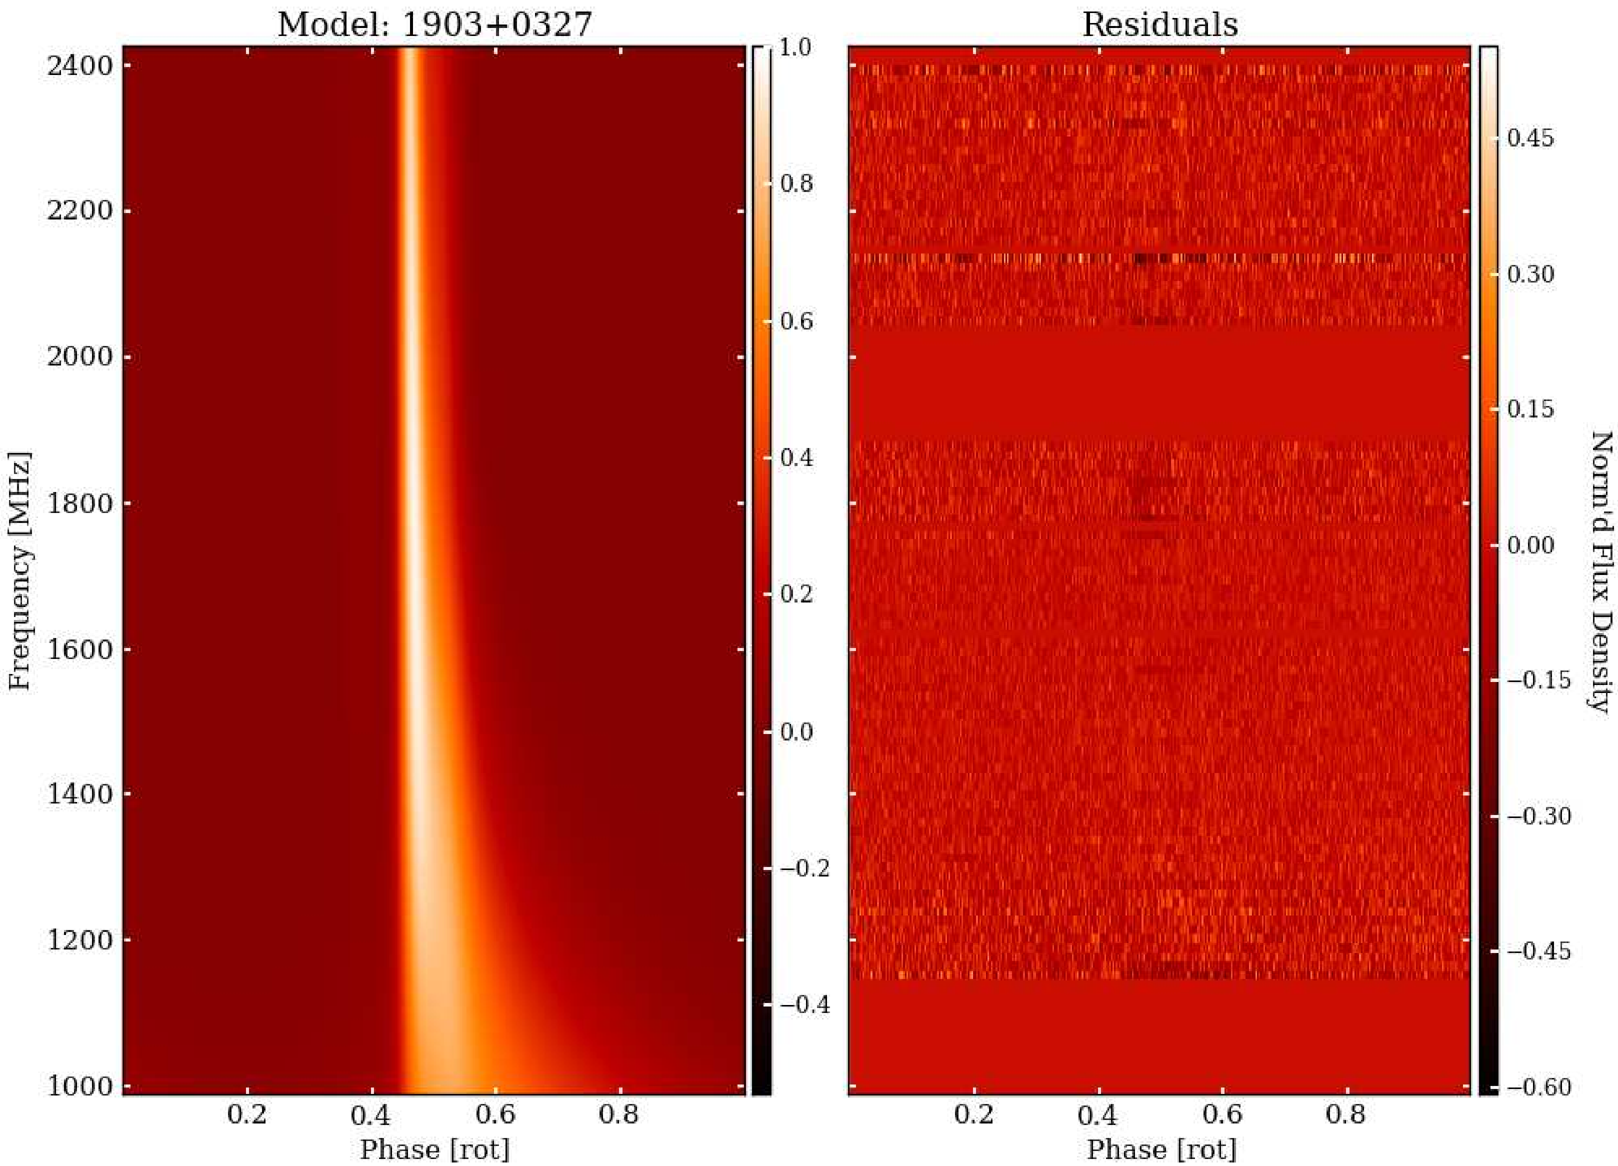

Fig. 4.199 – Fitted Gaussian component model and residuals for J1903+0327.

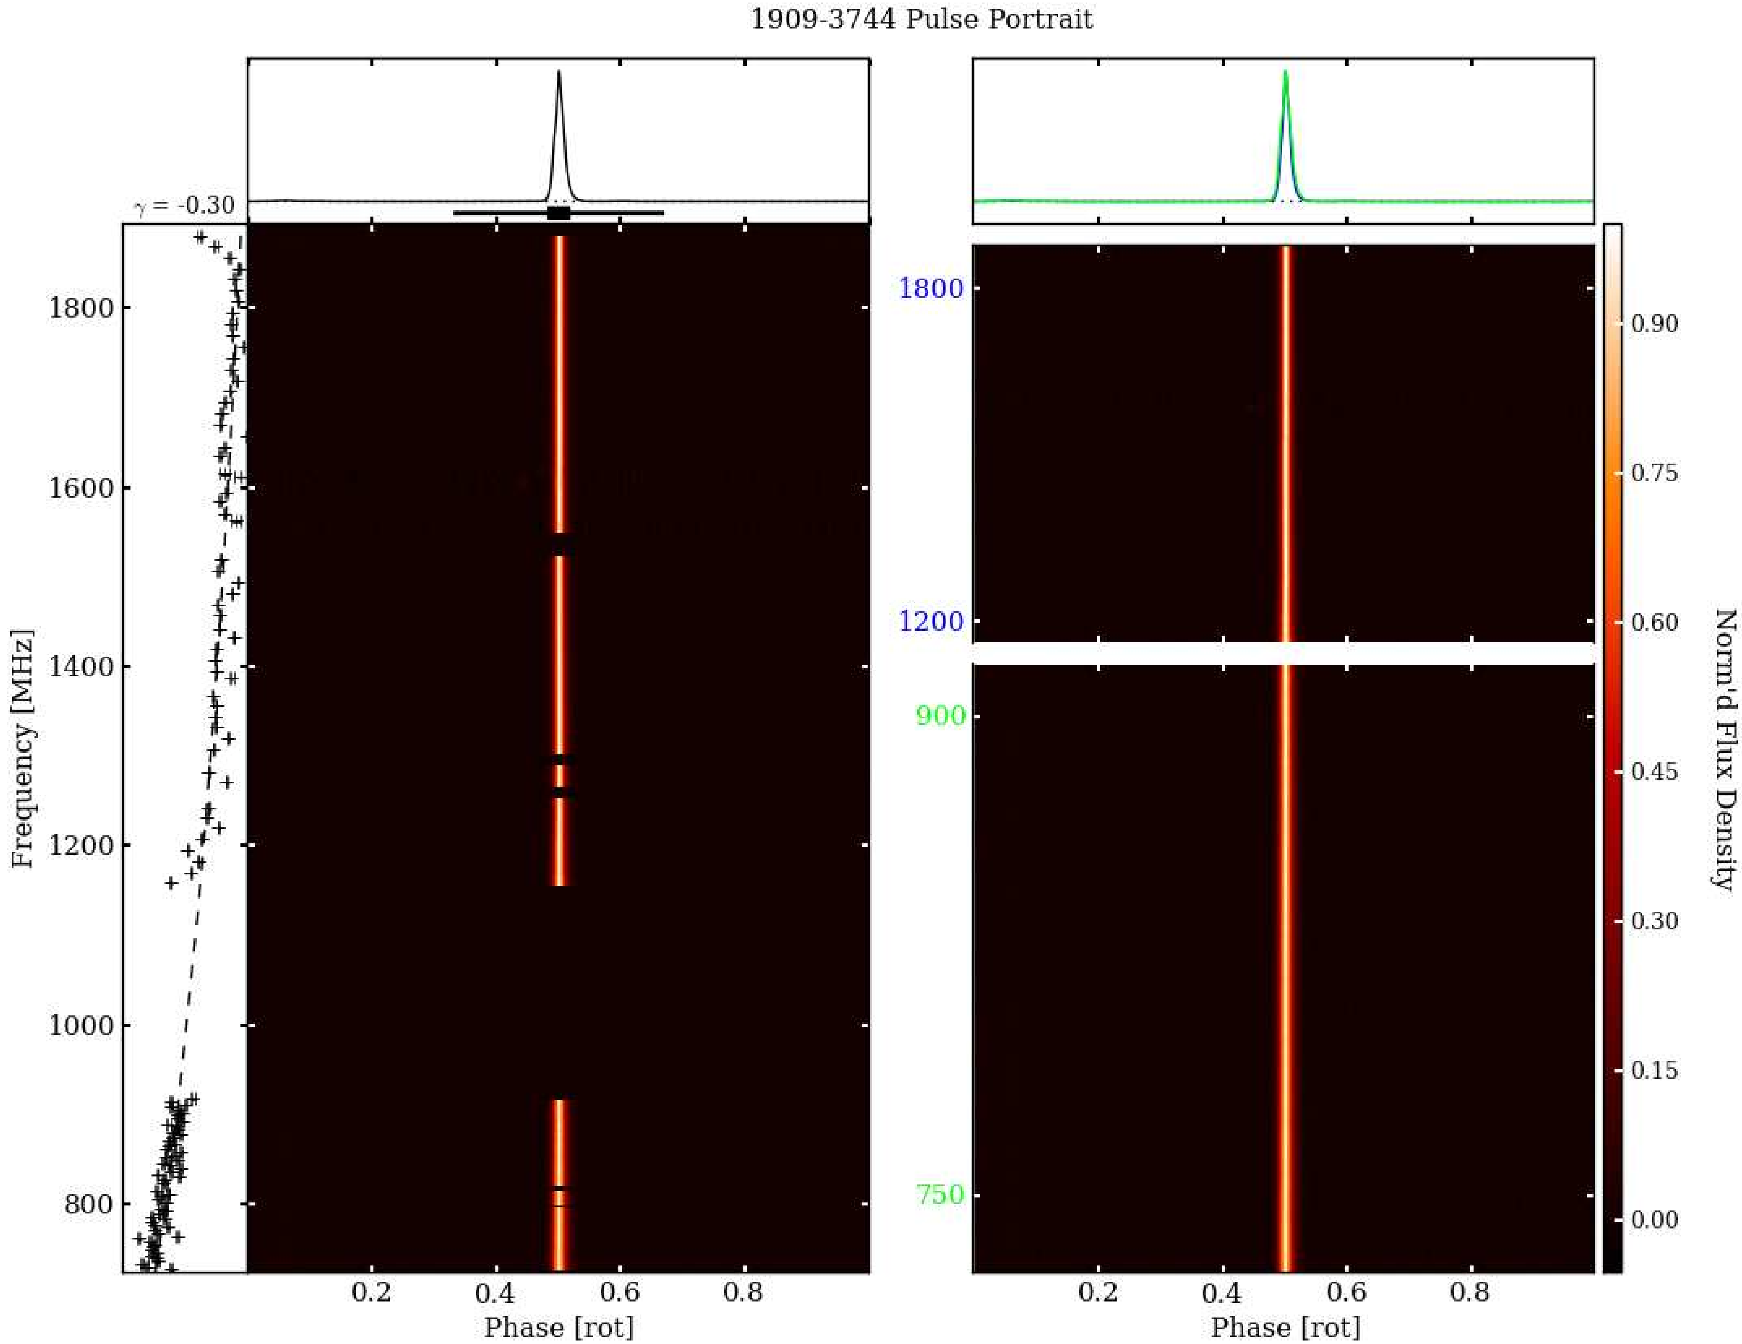

Fig. 4.200 – Concatenated portrait data and average profiles for J1909−3744.

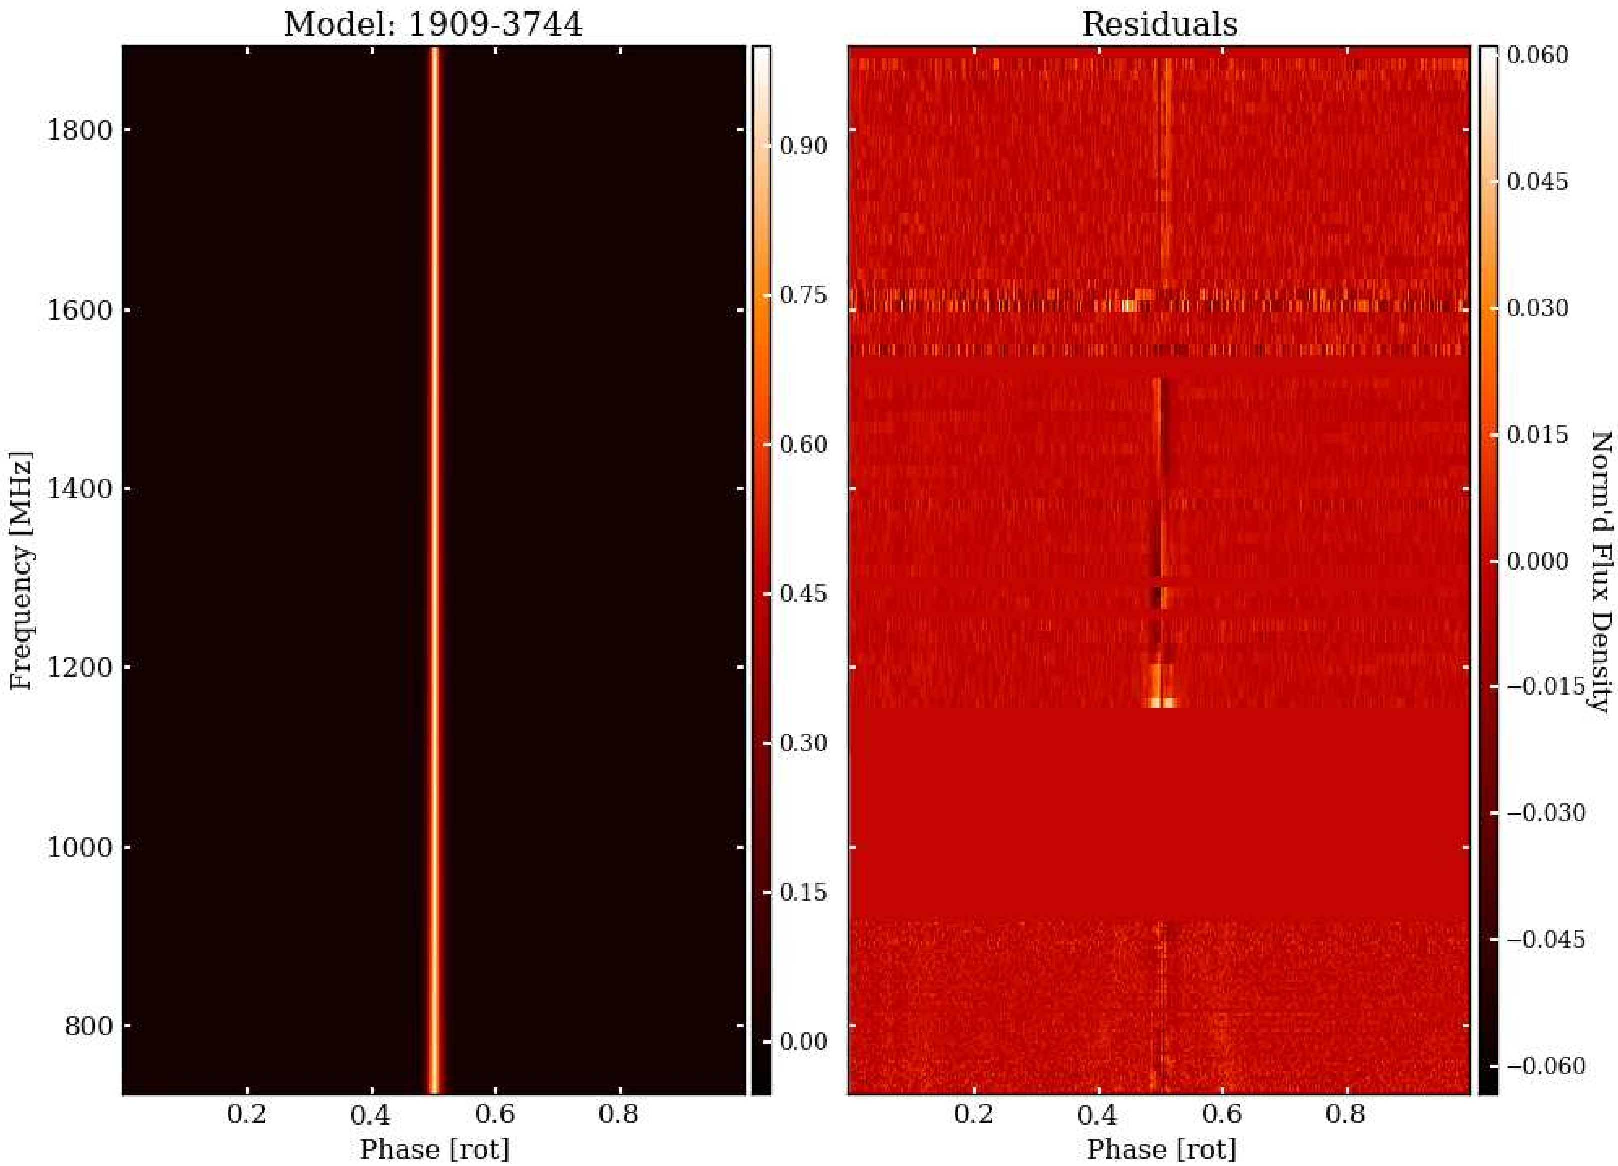

Fig. 4.201 – Fitted Gaussian component model and residuals for J1909−3744.

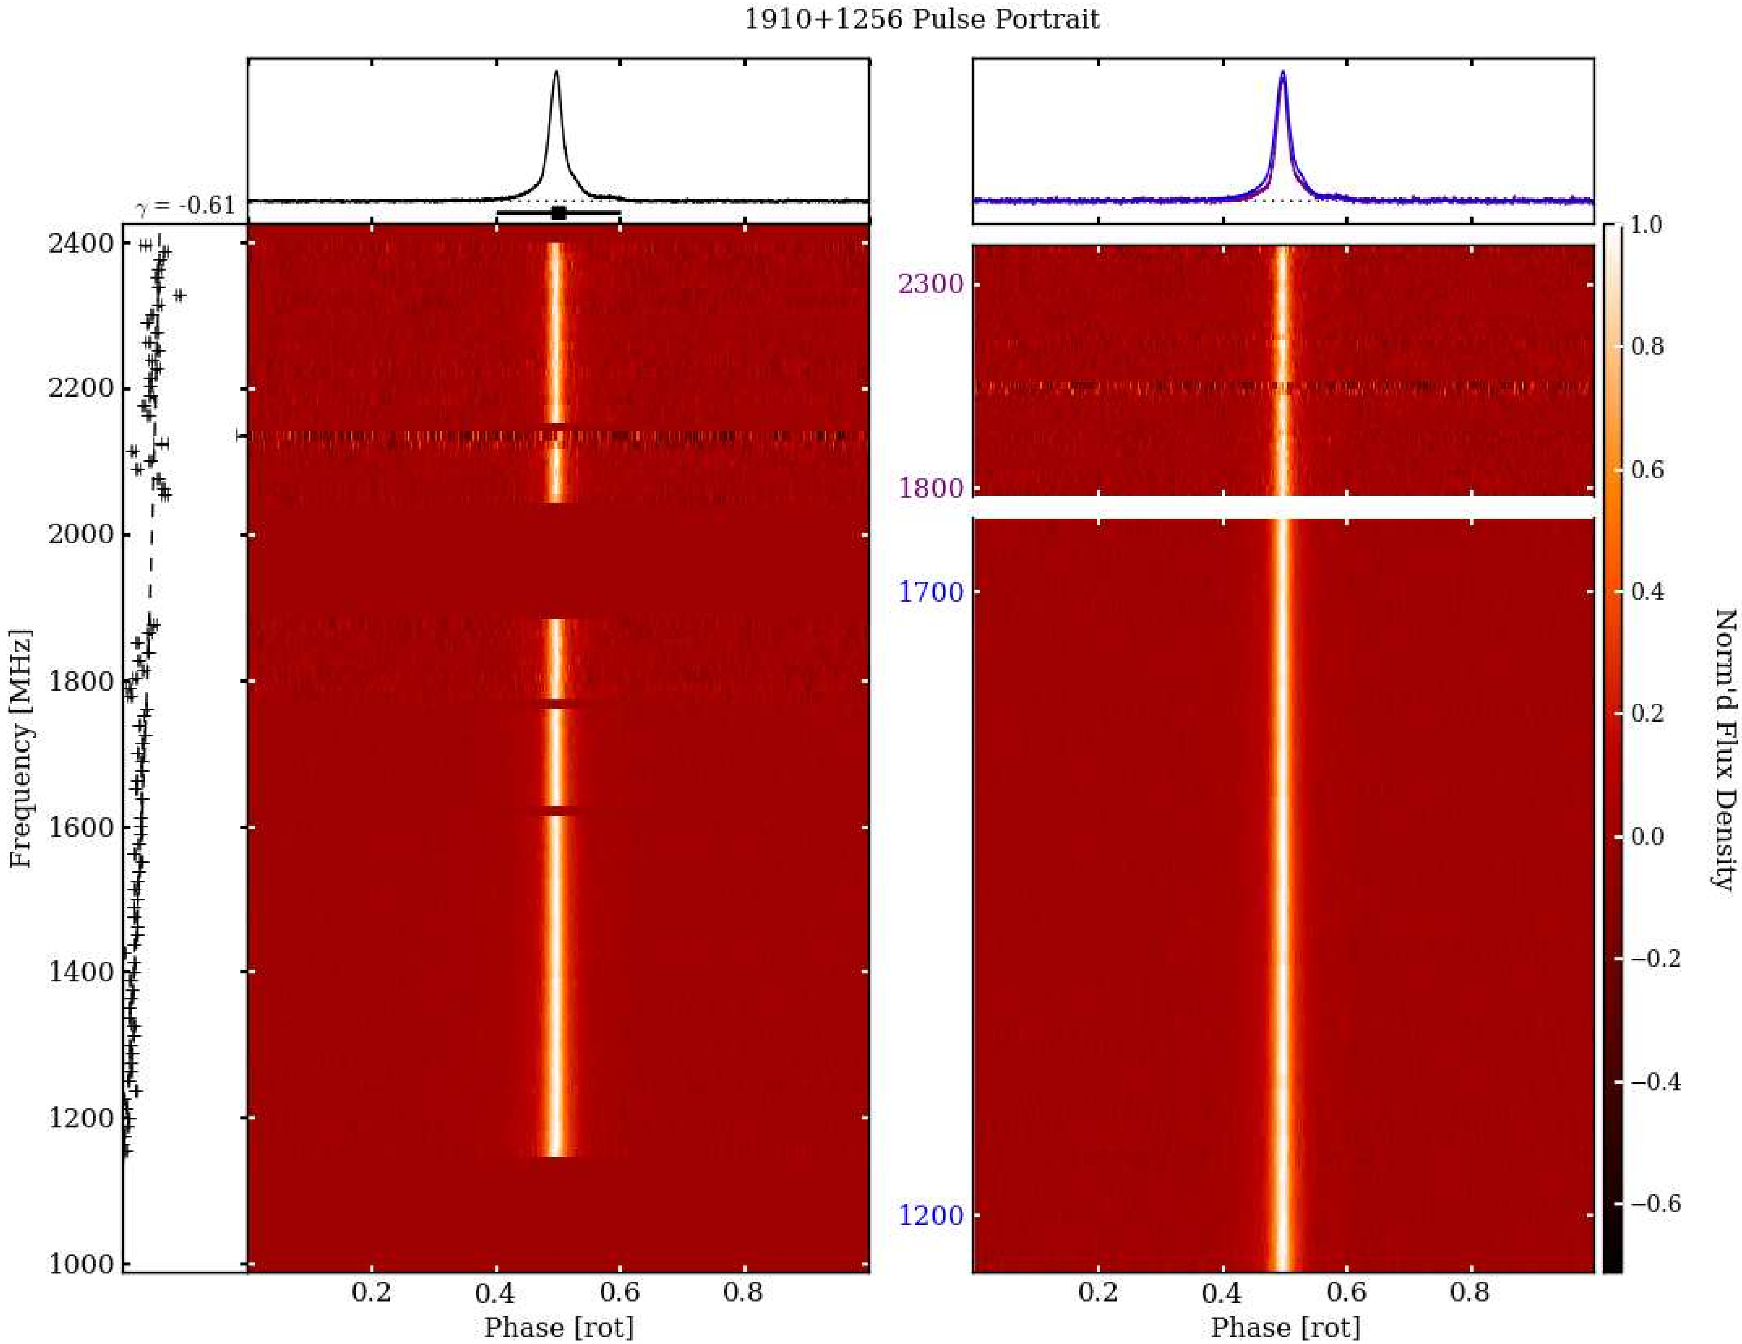

Fig. 4.202 – Concatenated portrait data and average profiles for J1910+1256.

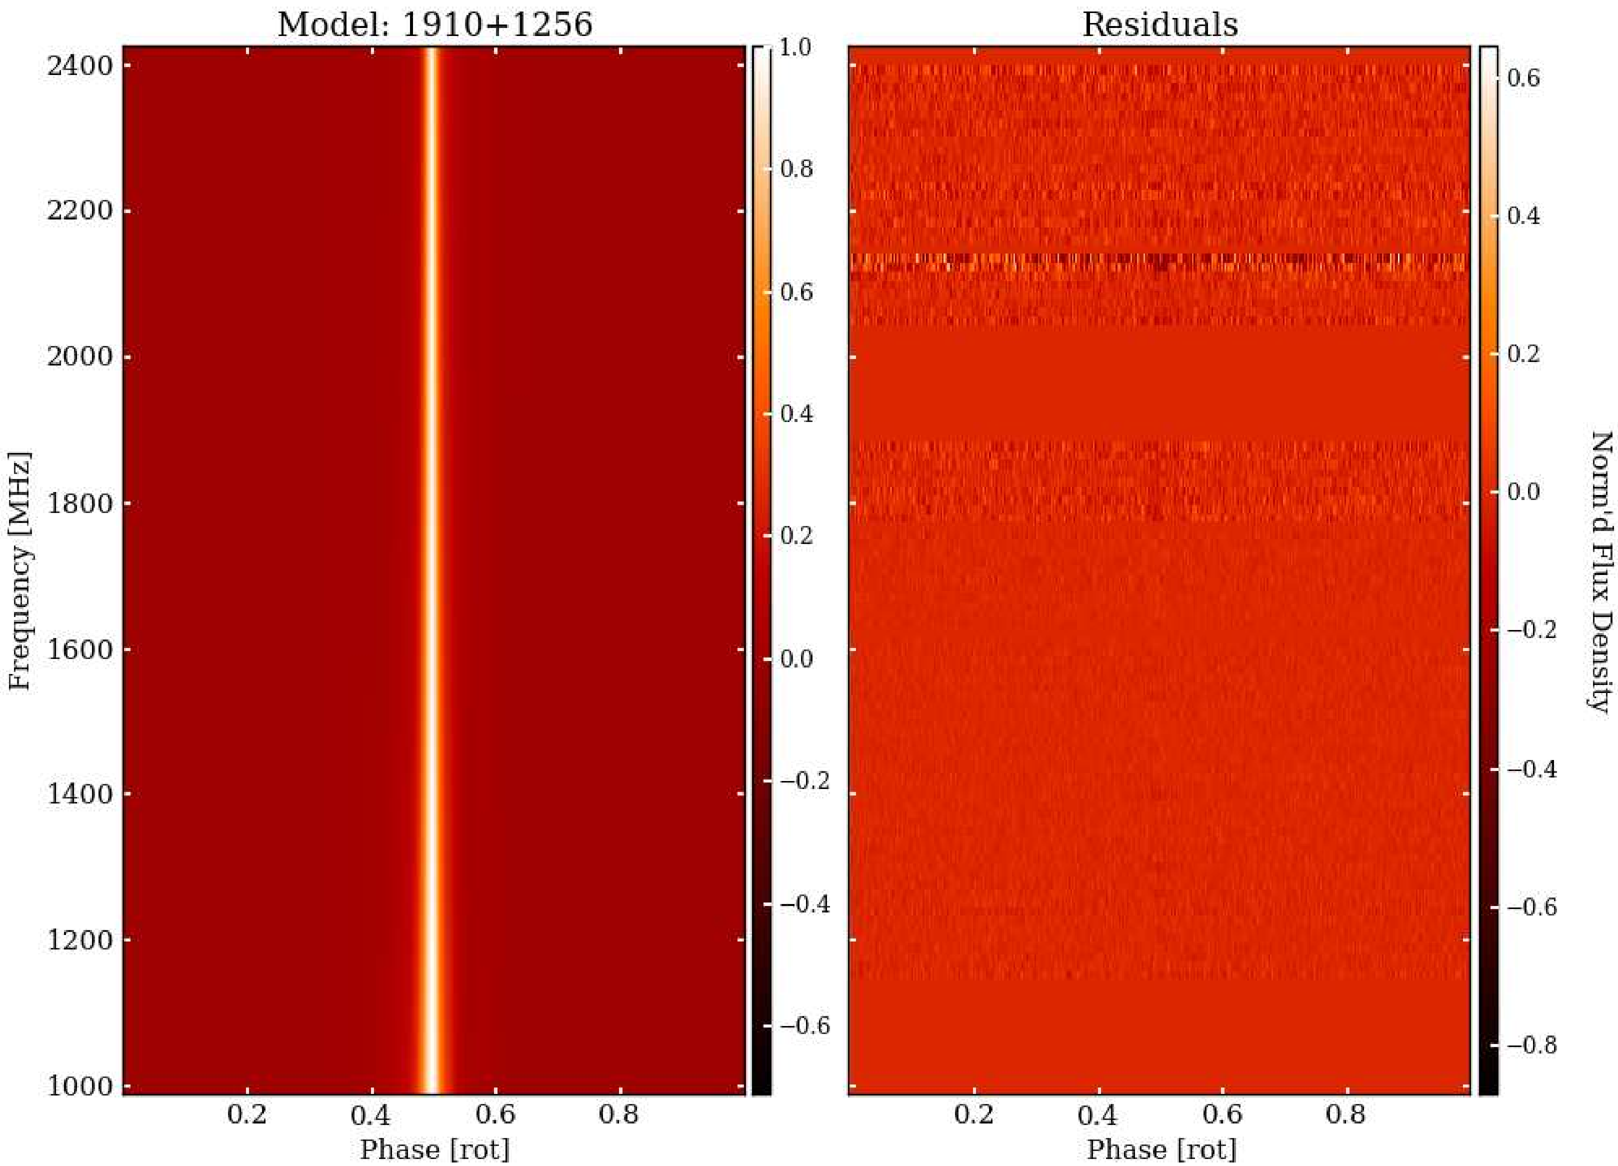

Fig. 4.203 – Fitted Gaussian component model and residuals for J1910+1256.

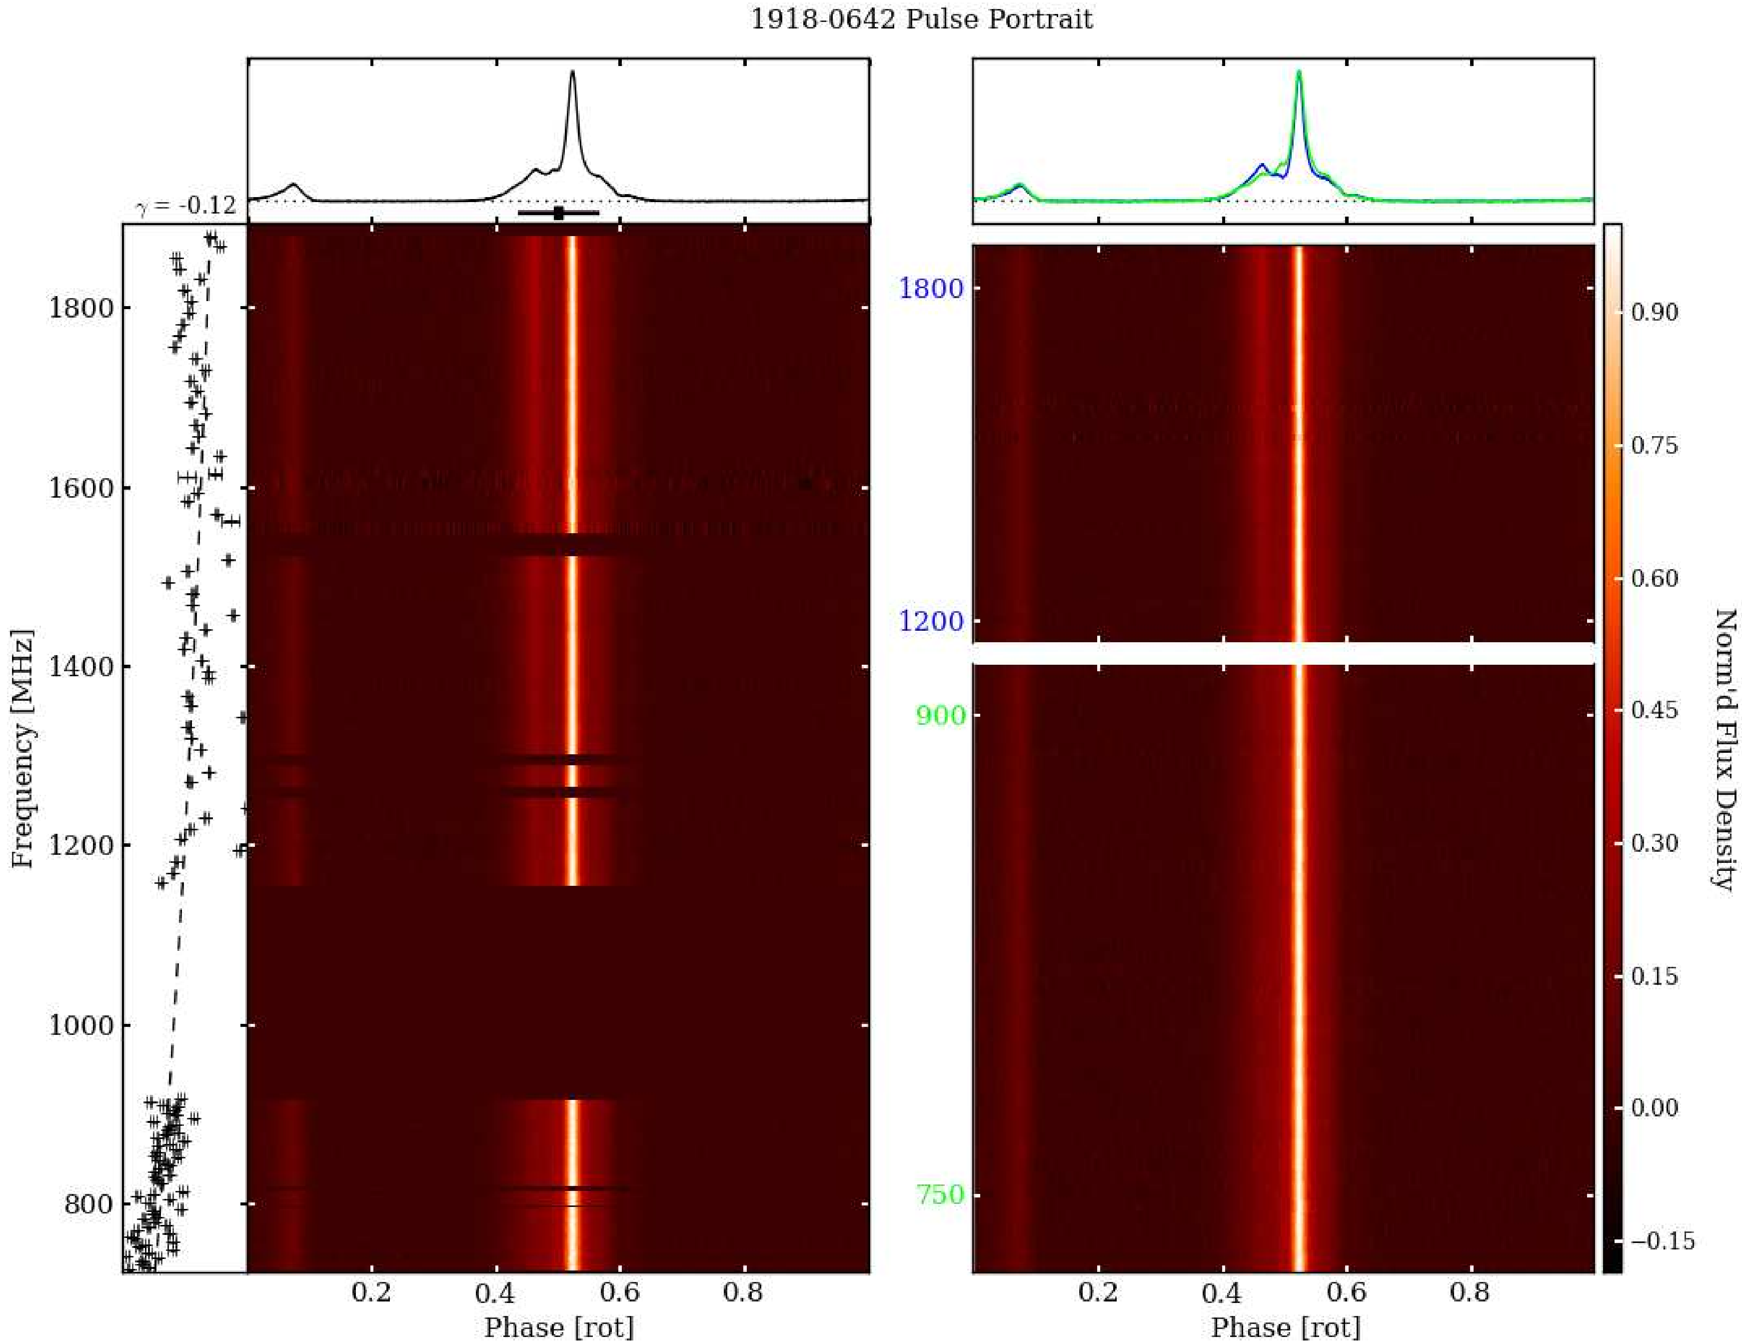

Fig. 4.204 – Concatenated portrait data and average profiles for J1918−0642.

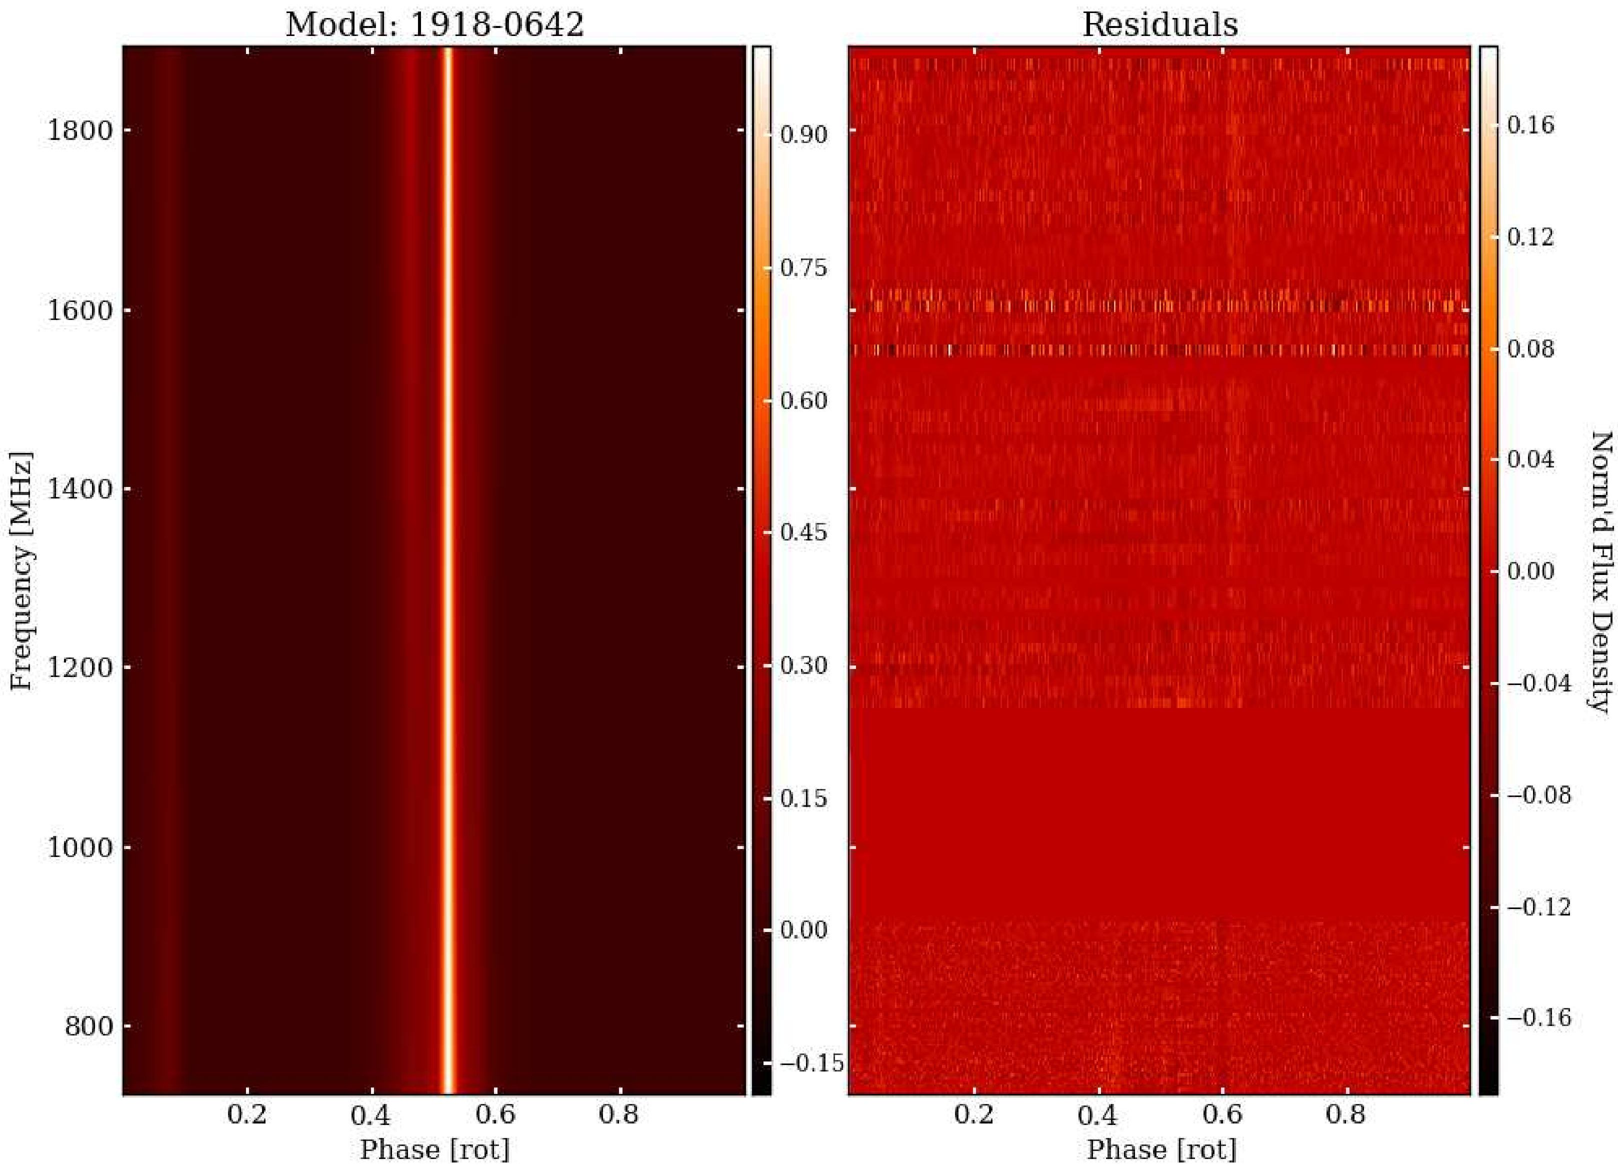

Fig. 4.205 – Fitted Gaussian component model and residuals for J1918−0642.

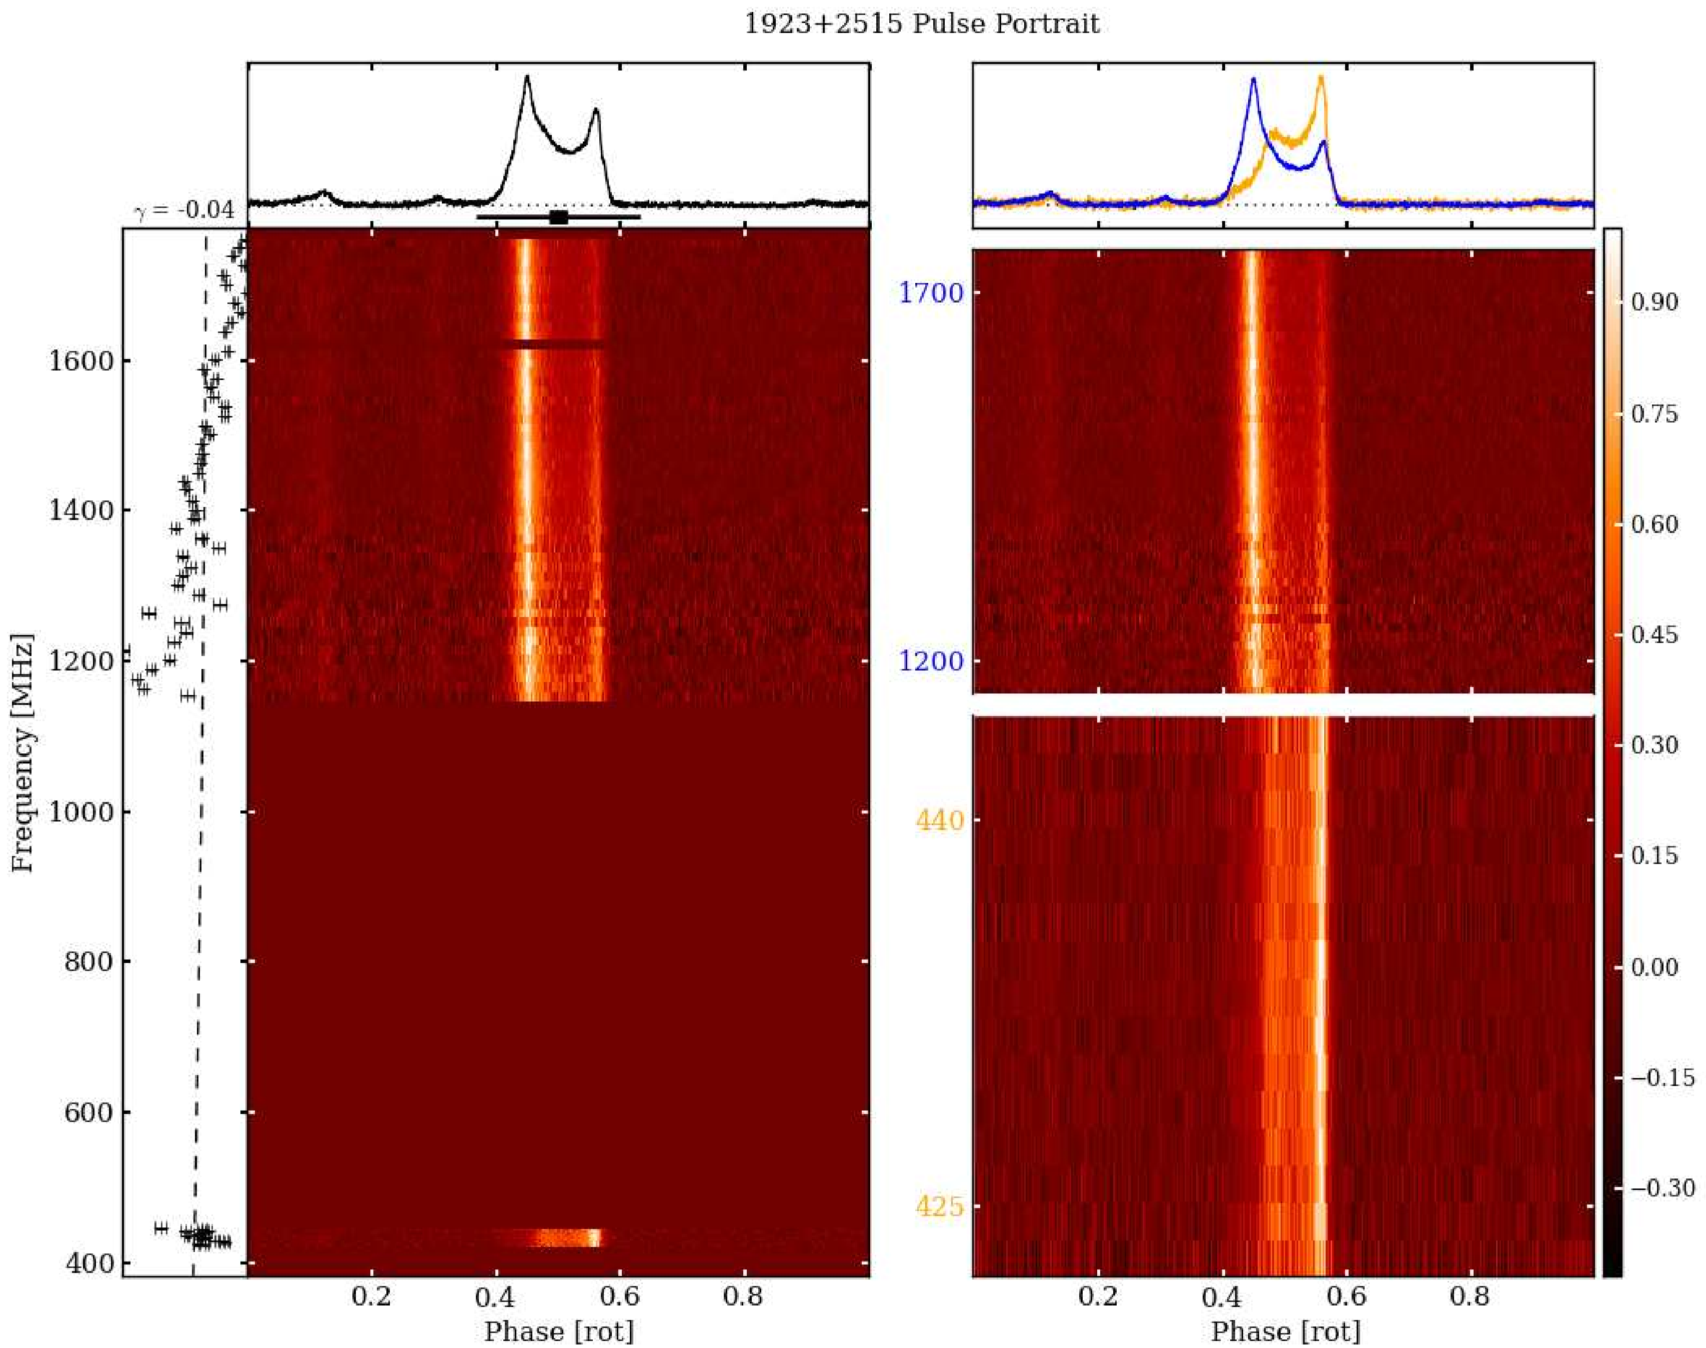

Fig. 4.206 – Concatenated portrait data and average profiles for J1923+2515.

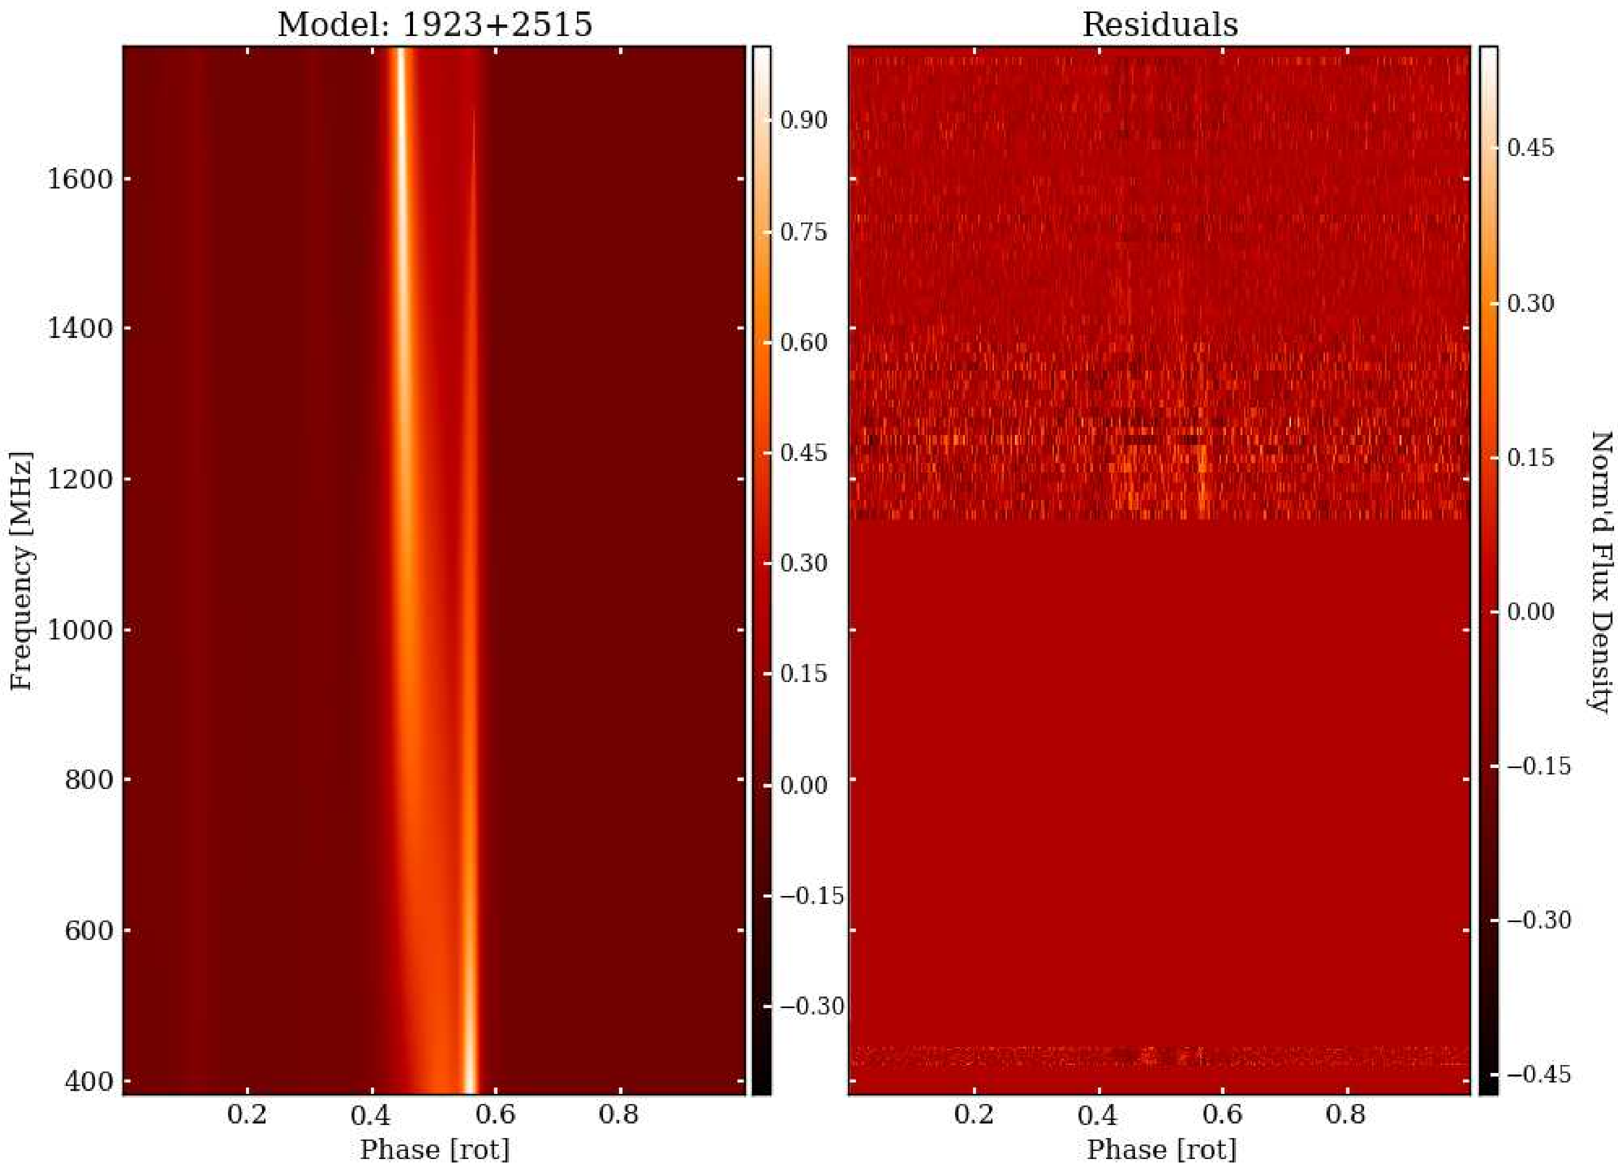

Fig. 4.207 – Fitted Gaussian component model and residuals for J1923+2515.

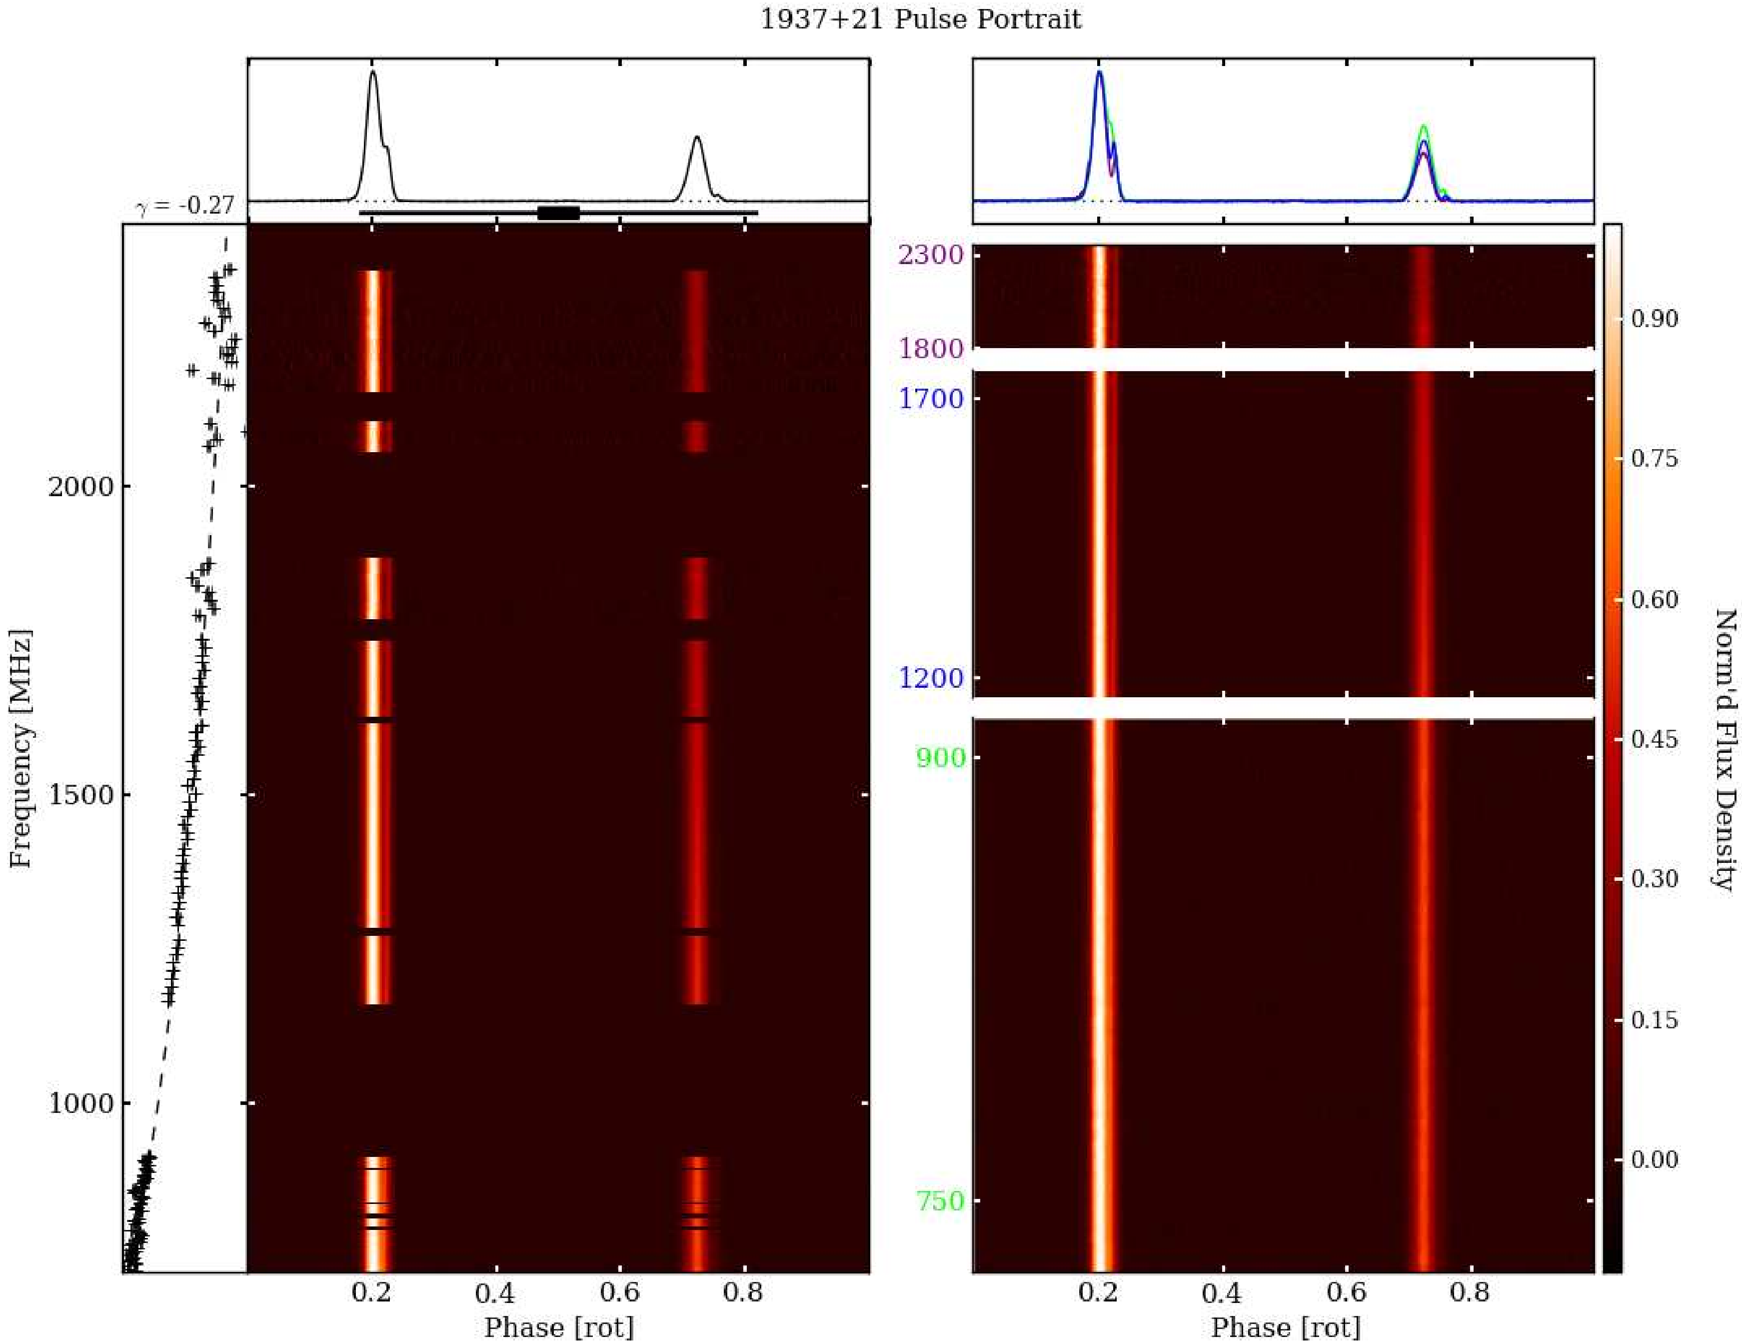

Fig. 4.208 – Concatenated portrait data and average profiles for B1937+21.

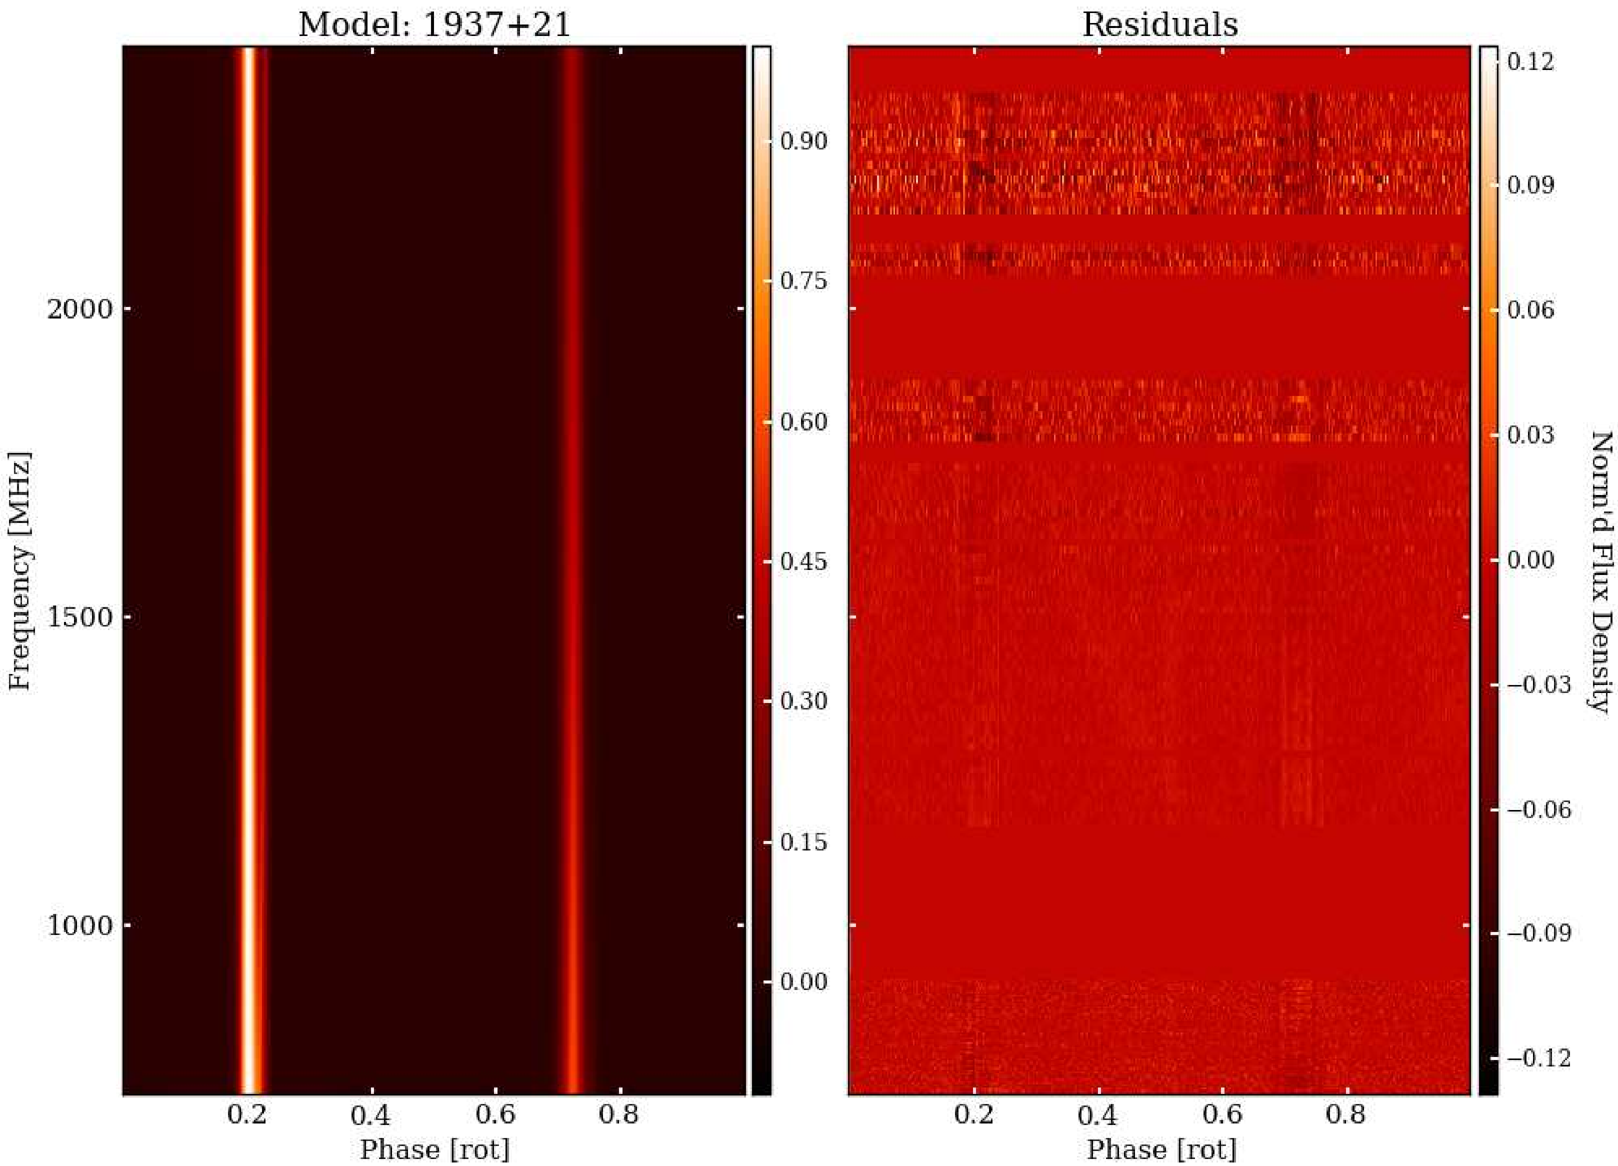

Fig. 4.209 – Fitted Gaussian component model and residuals for B1937+21.

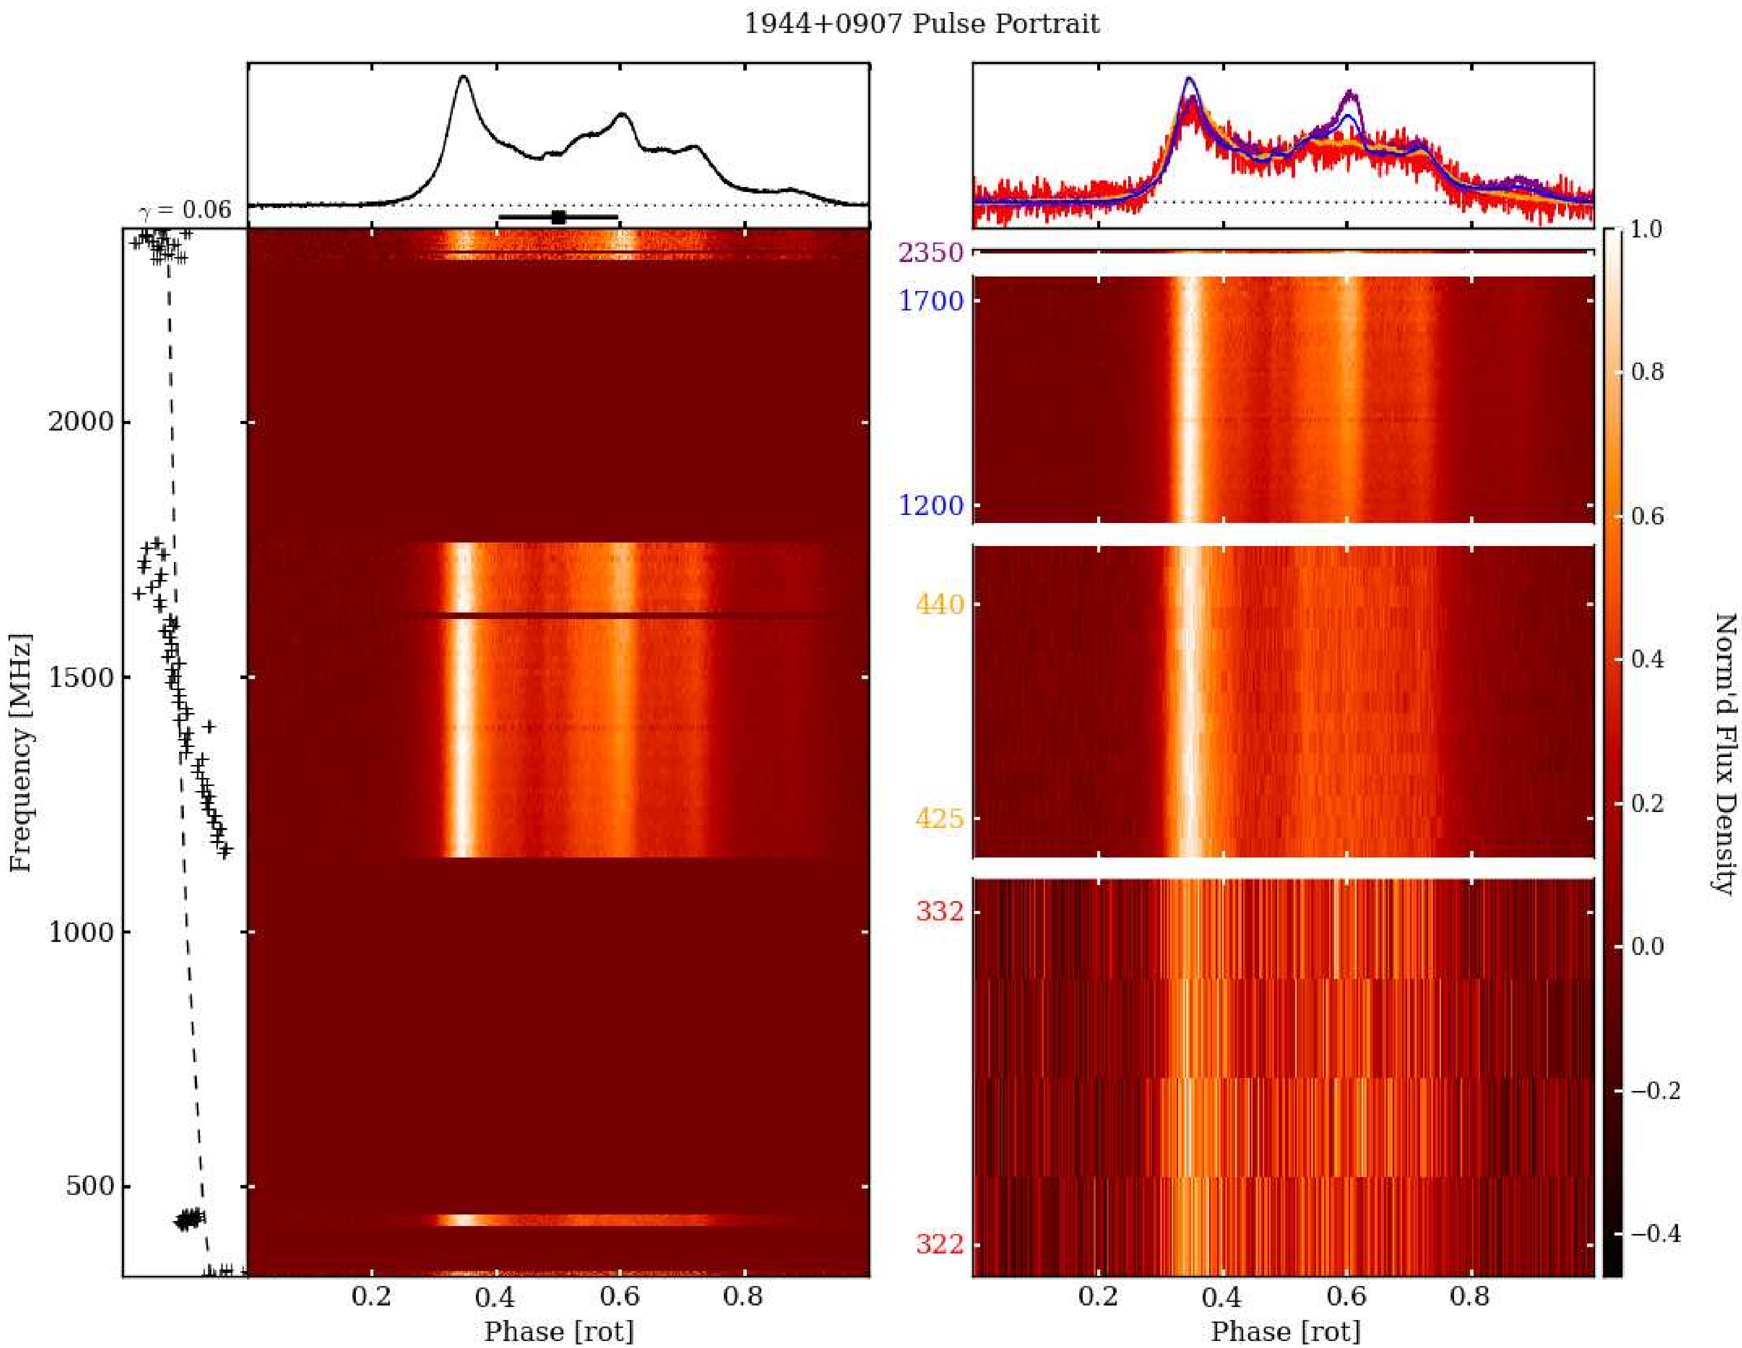

Fig. 4.210 – Concatenated portrait data and average profiles for J1944+0907.

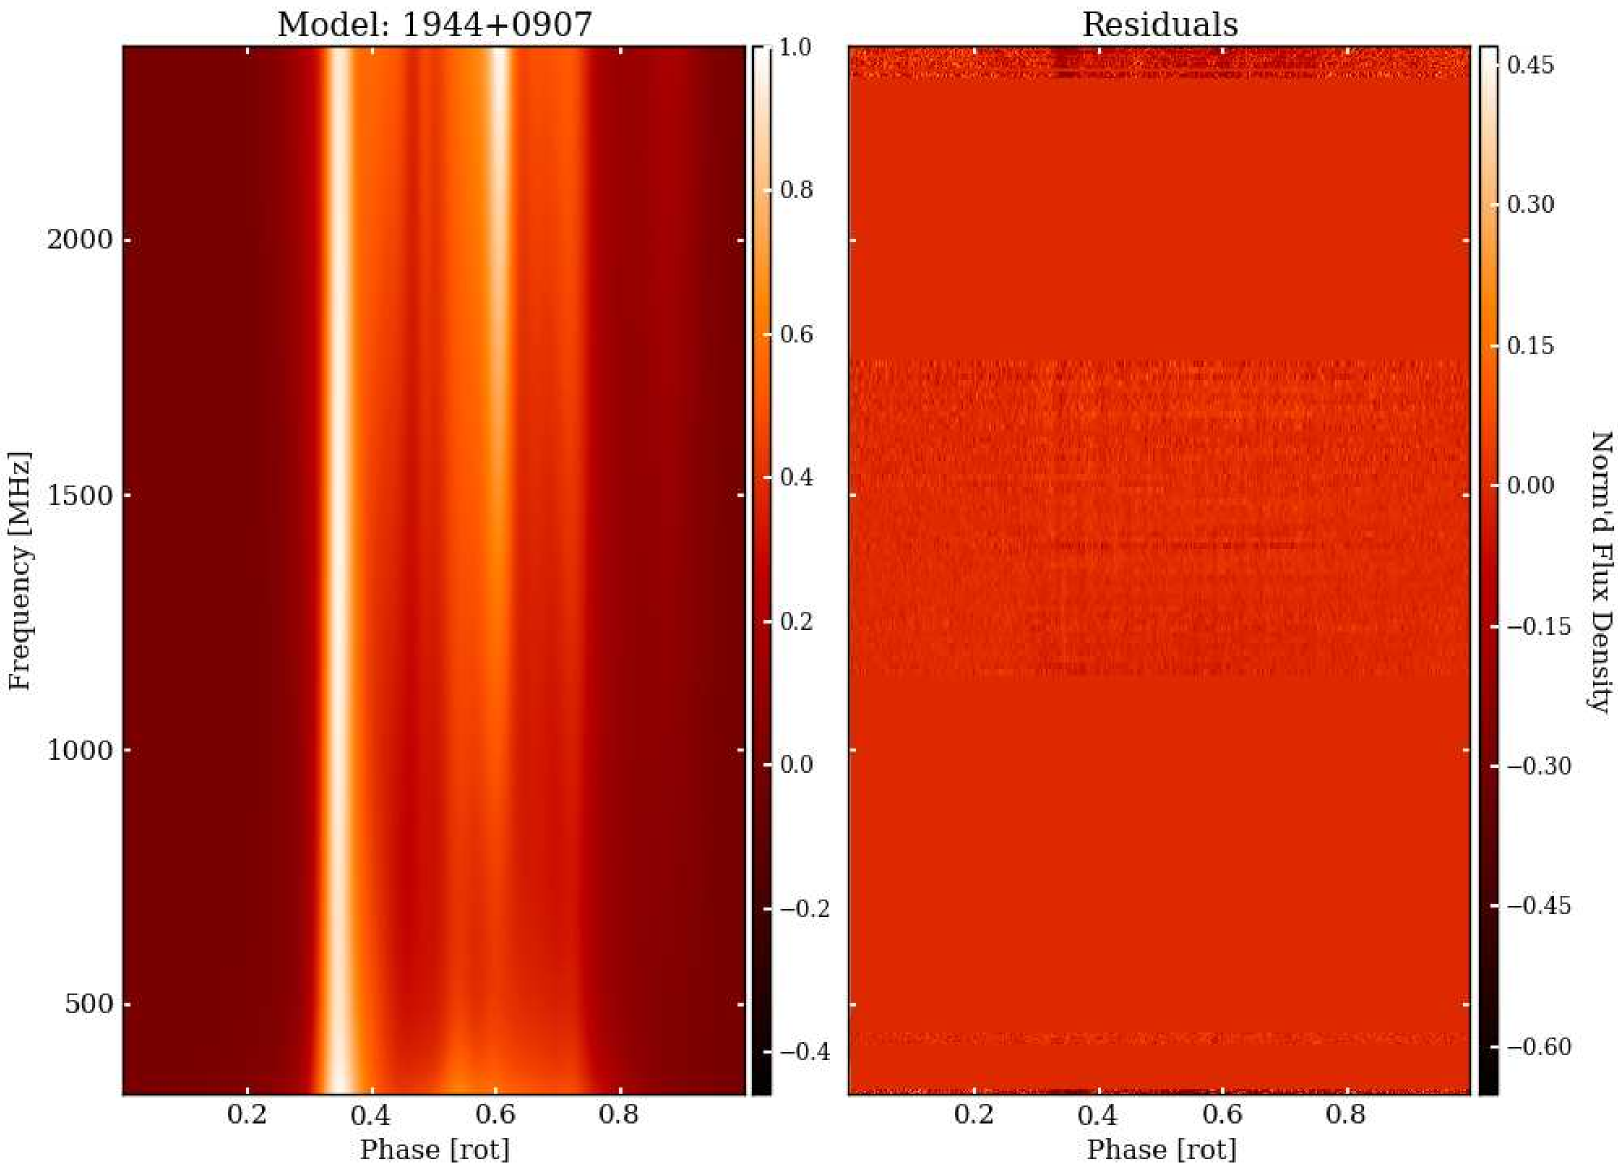

Fig. 4.211 – Fitted Gaussian component model and residuals for J1944+0907.

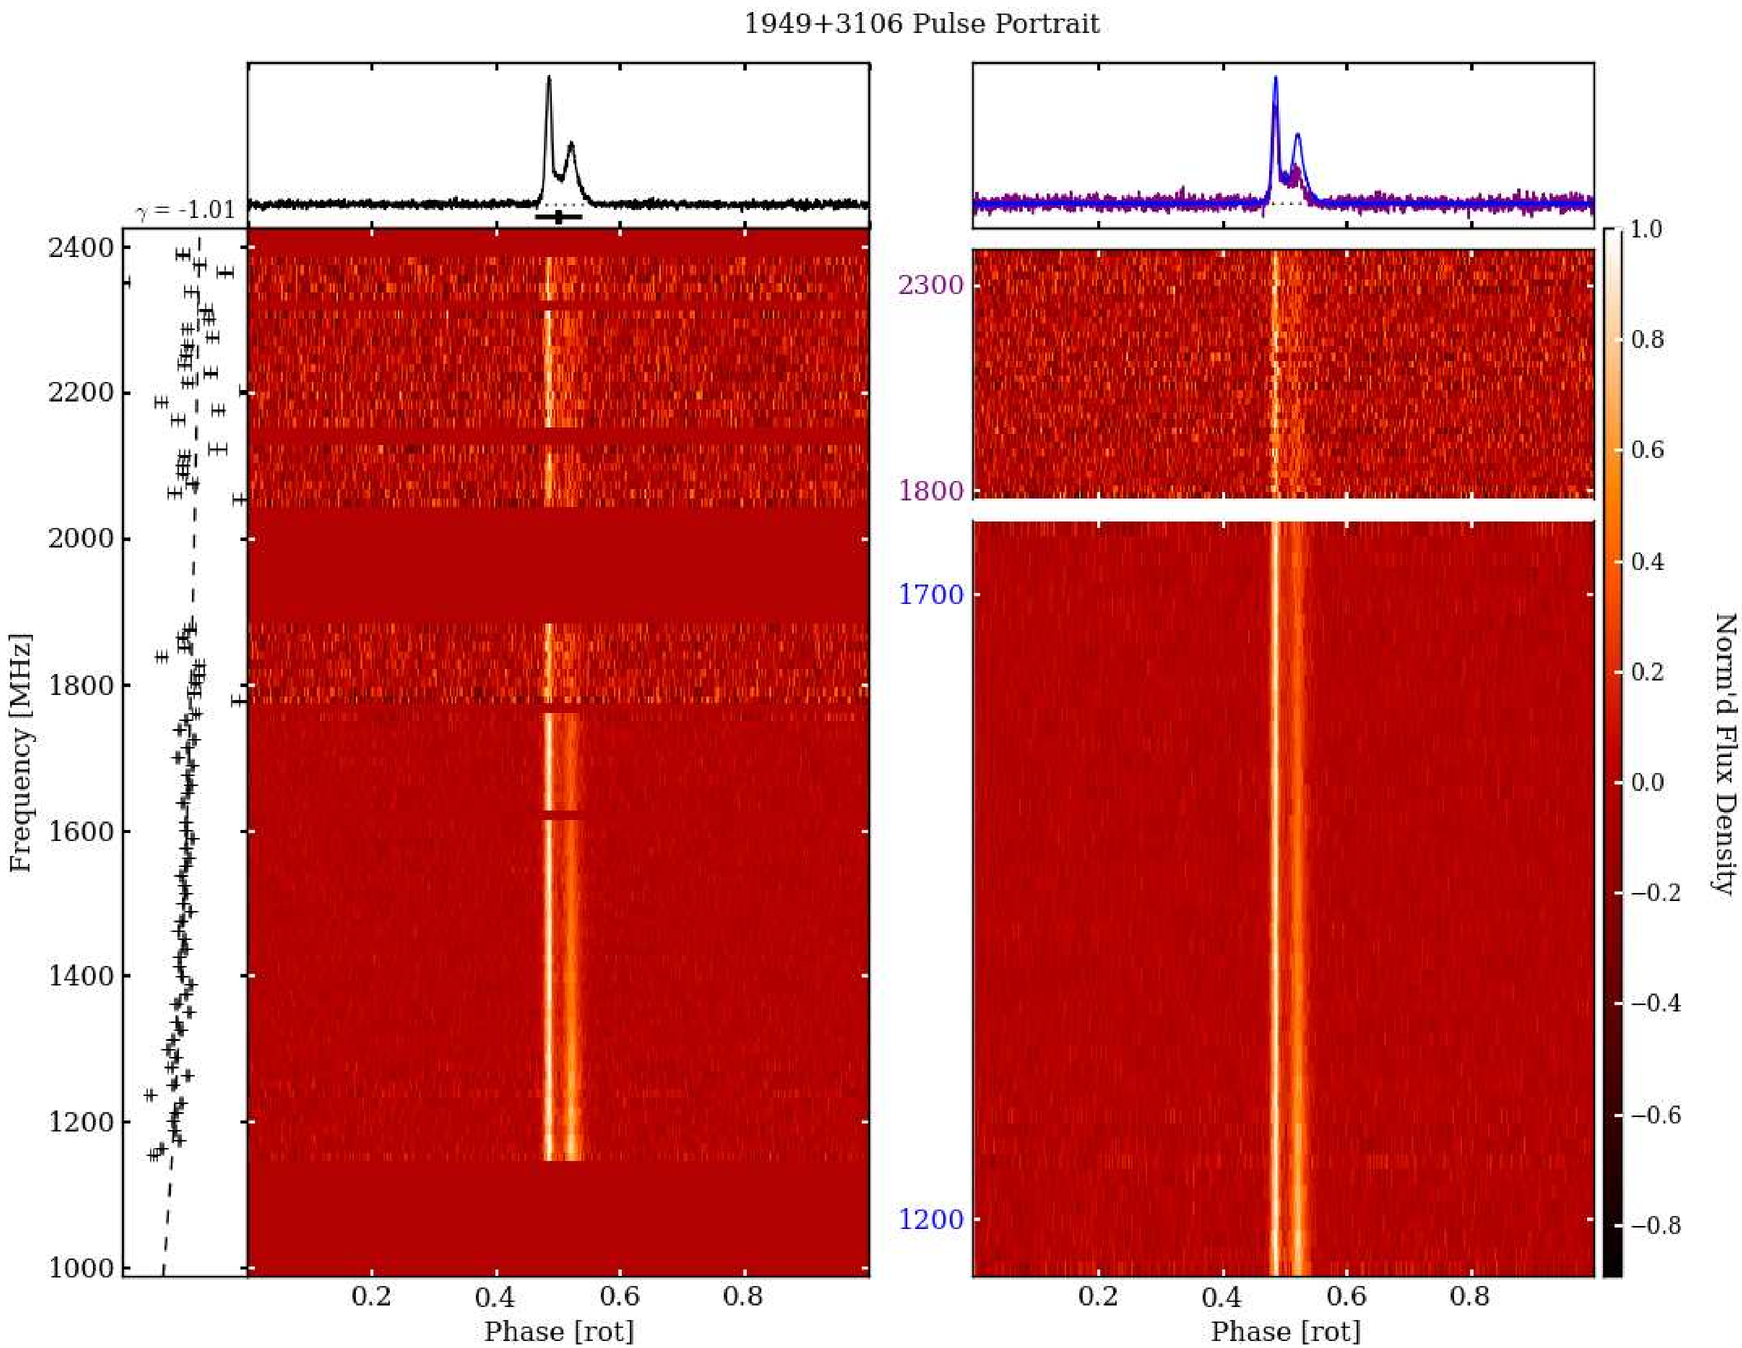

Fig. 4.212 – Concatenated portrait data and average profiles for J1949+3106.

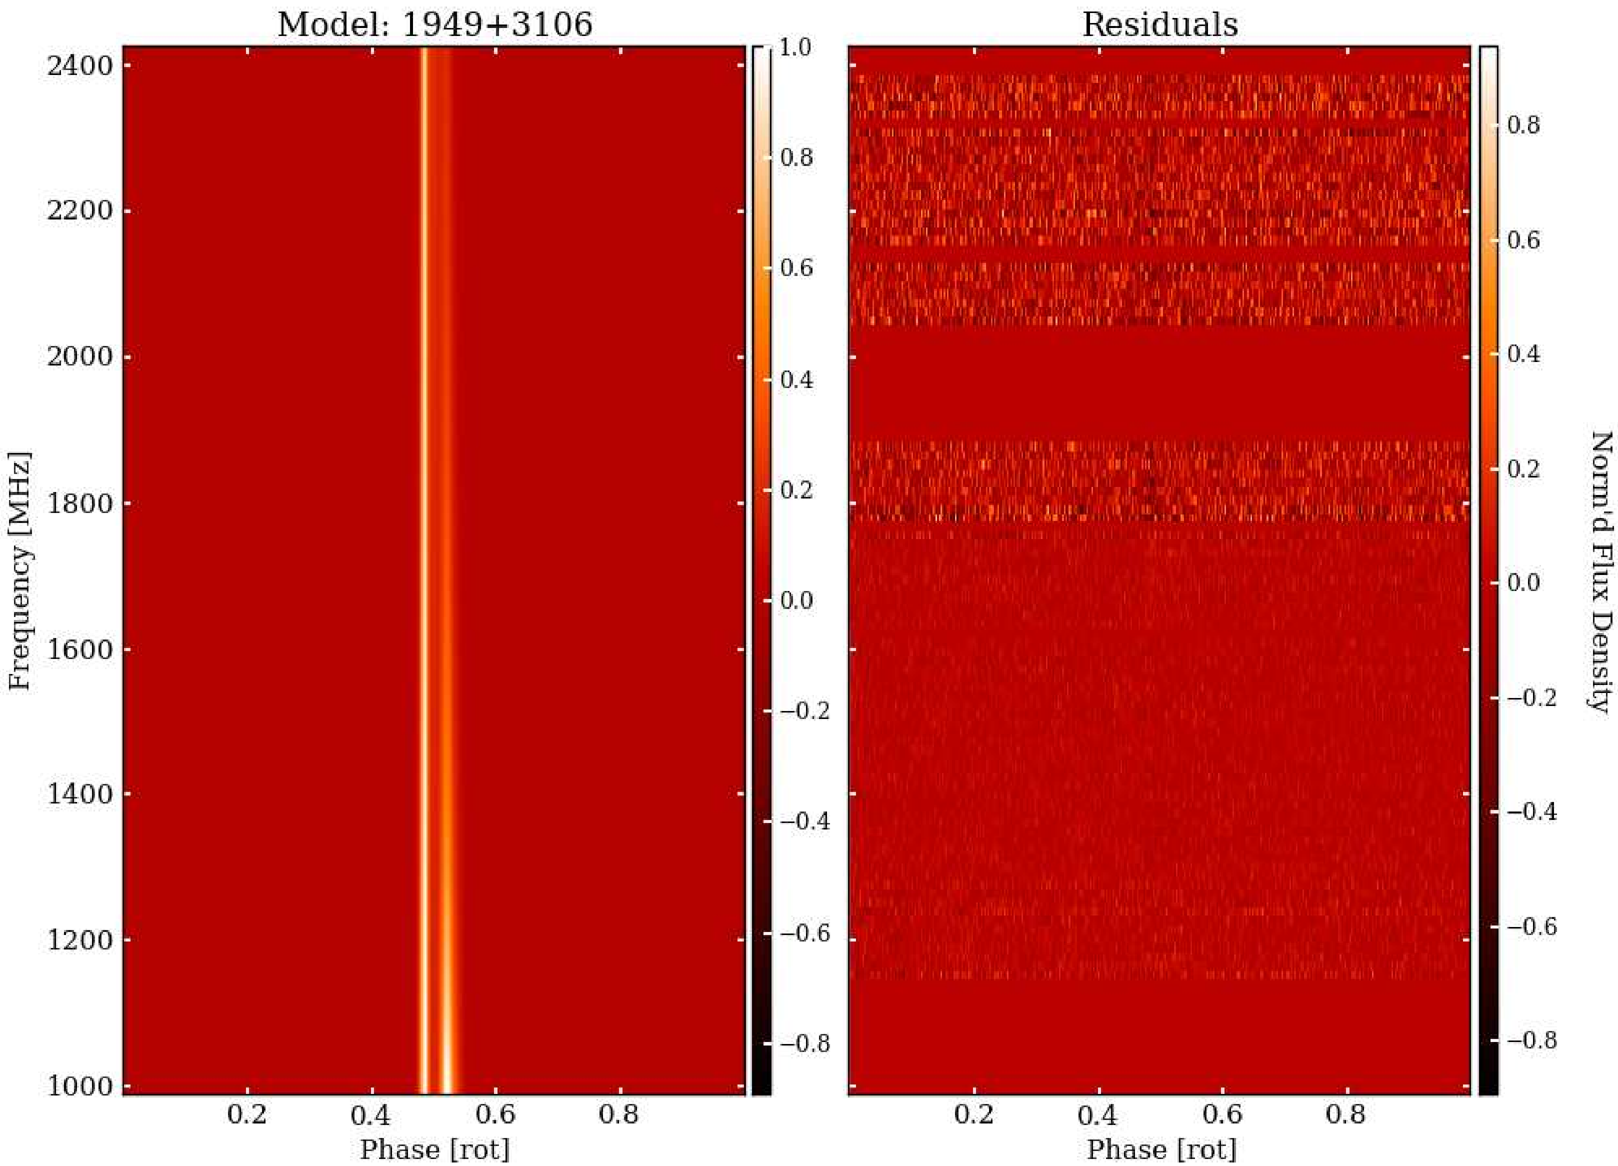

Fig. 4.213 – Fitted Gaussian component model and residuals for J1949+3106.

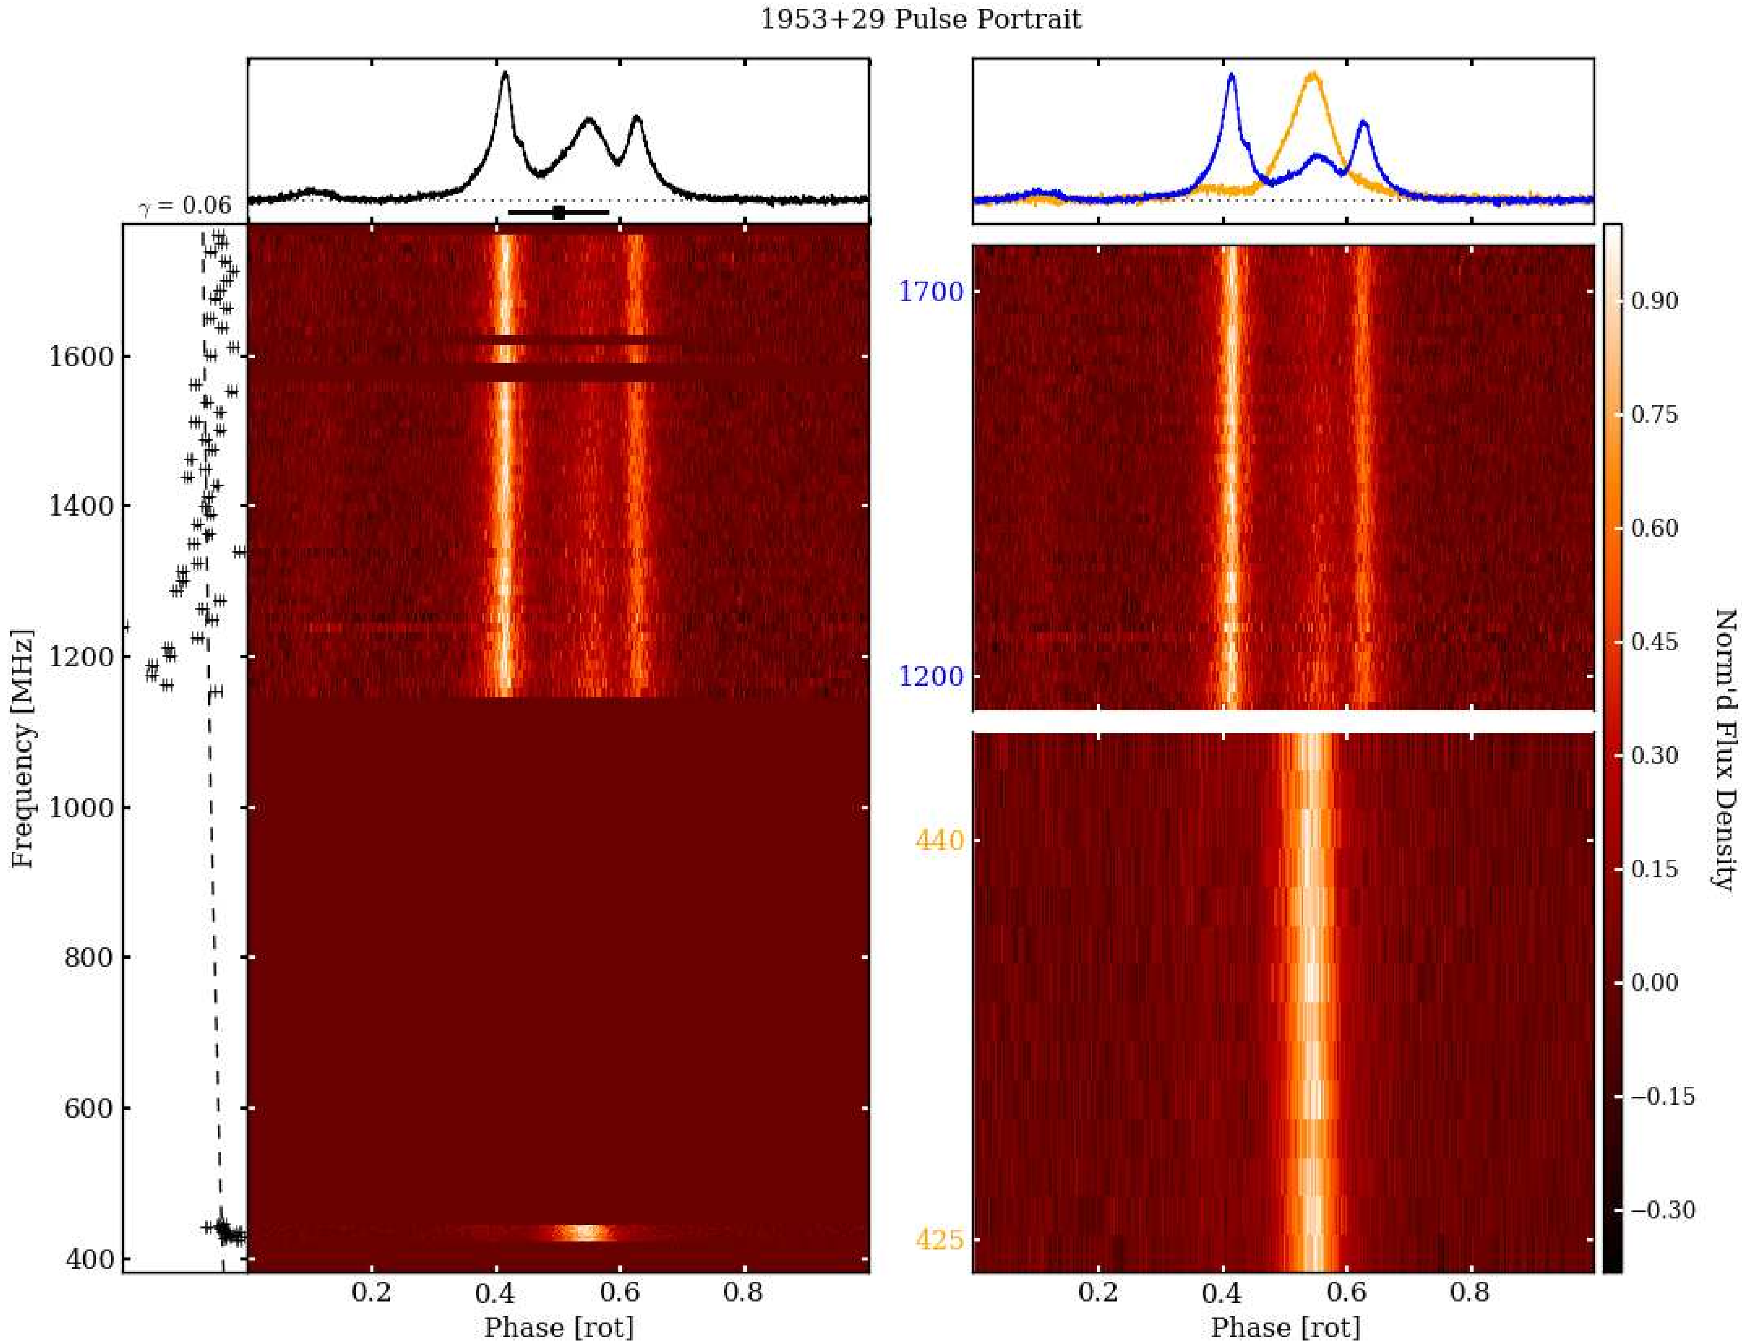

Fig. 4.214 – Concatenated portrait data and average profiles for B1953+29.

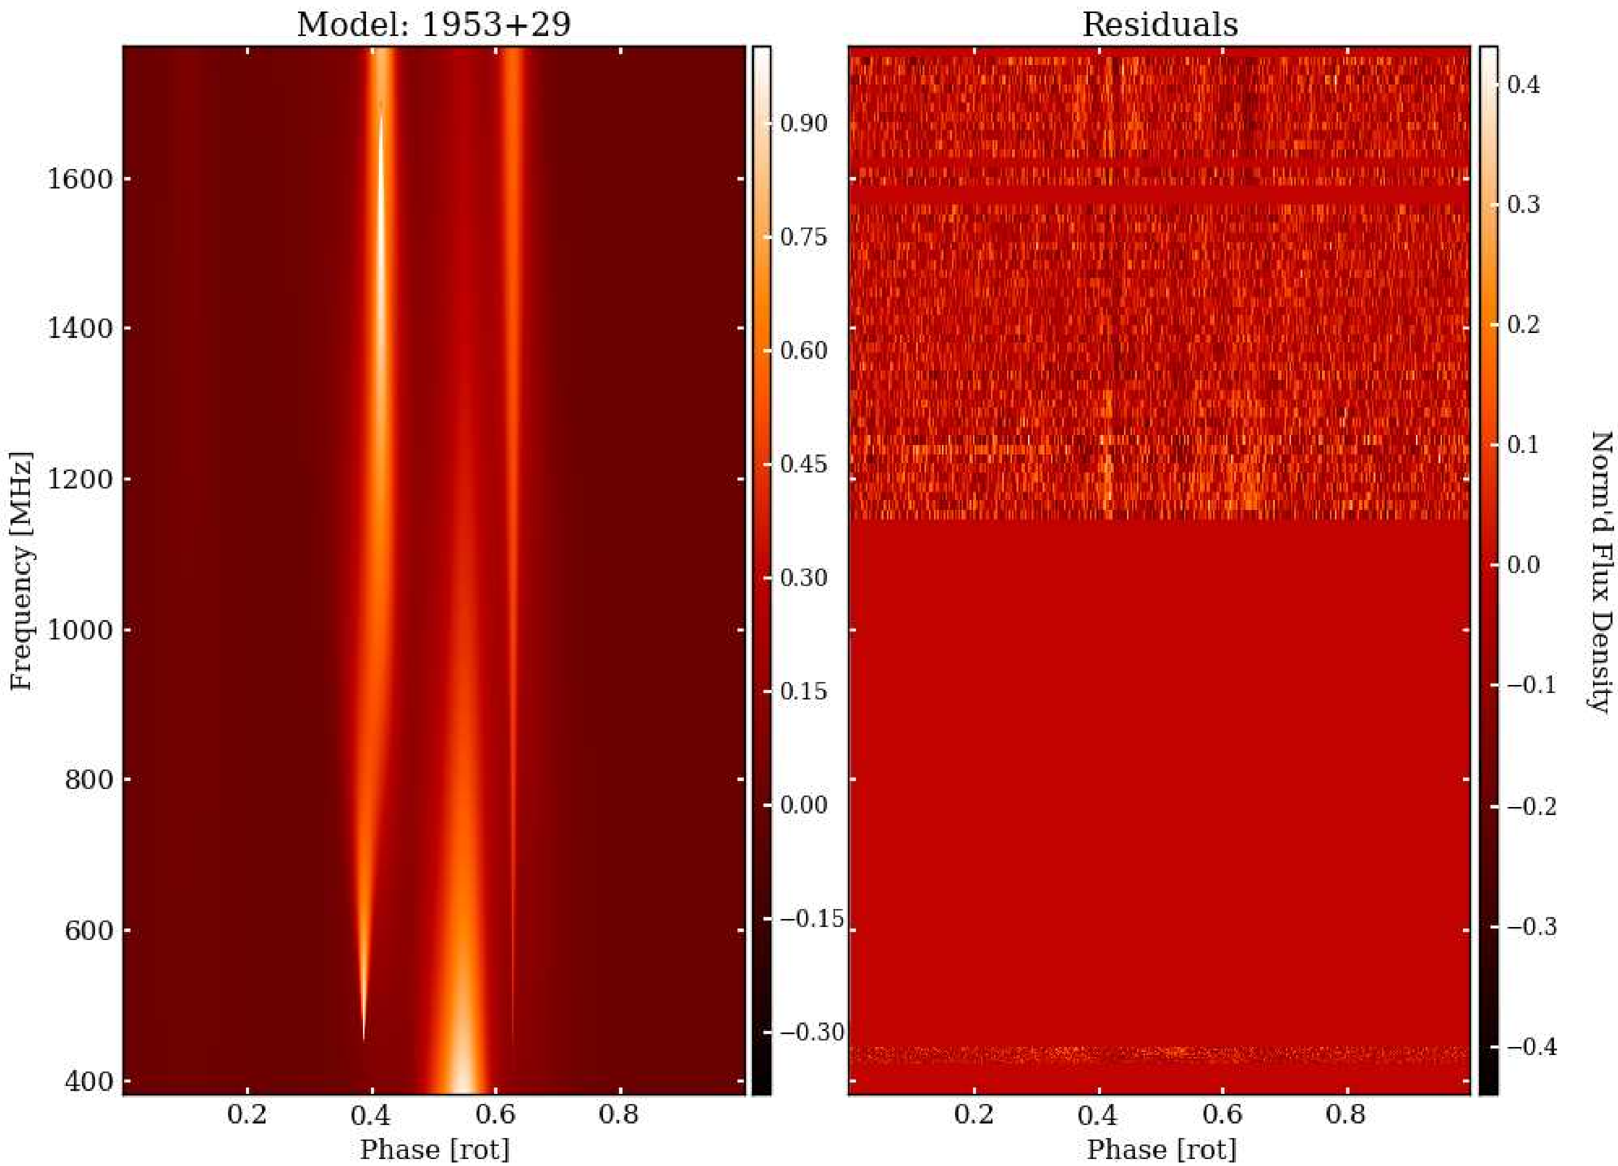

Fig. 4.215 – Fitted Gaussian component model and residuals for B1953+29.

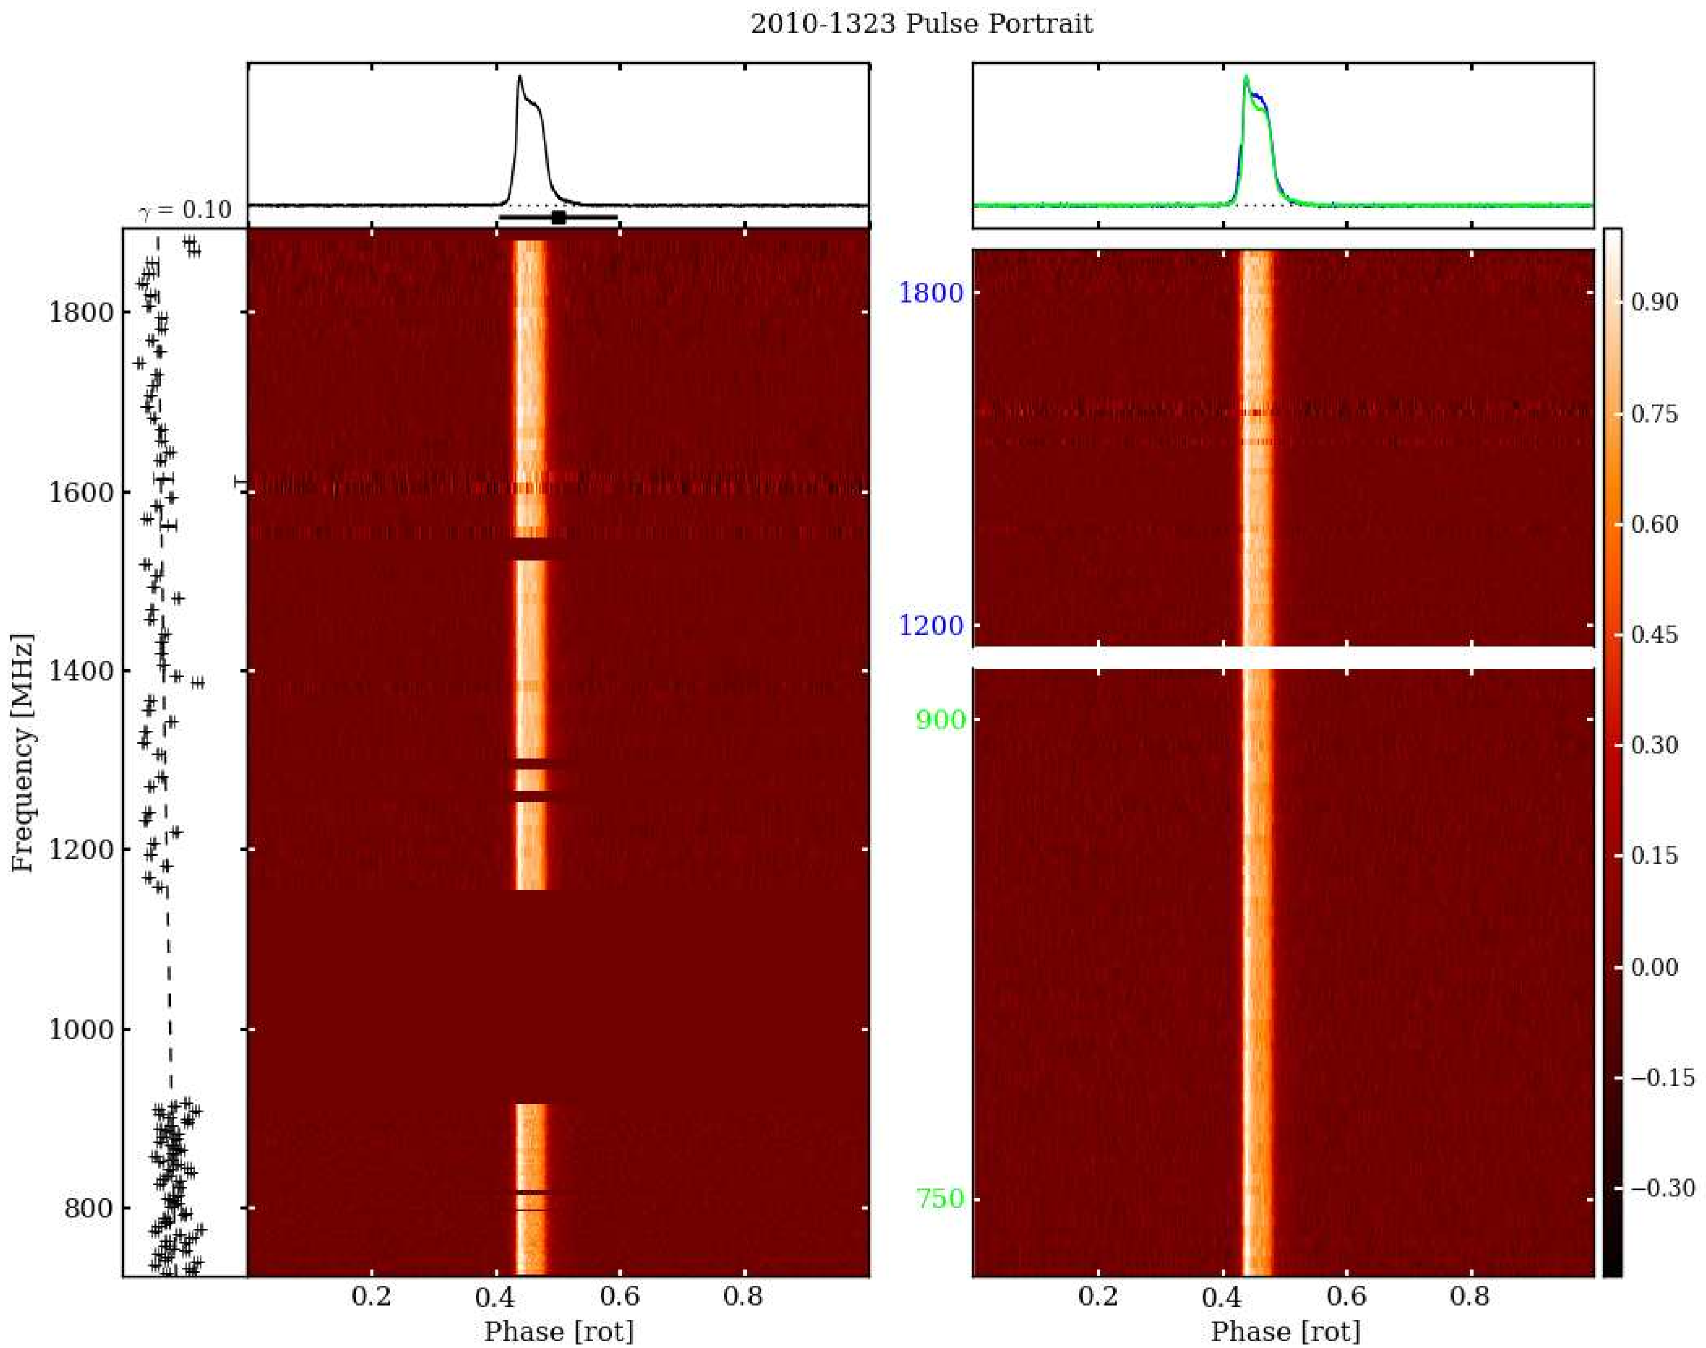

Fig. 4.216 – Concatenated portrait data and average profiles for J2010−1323.

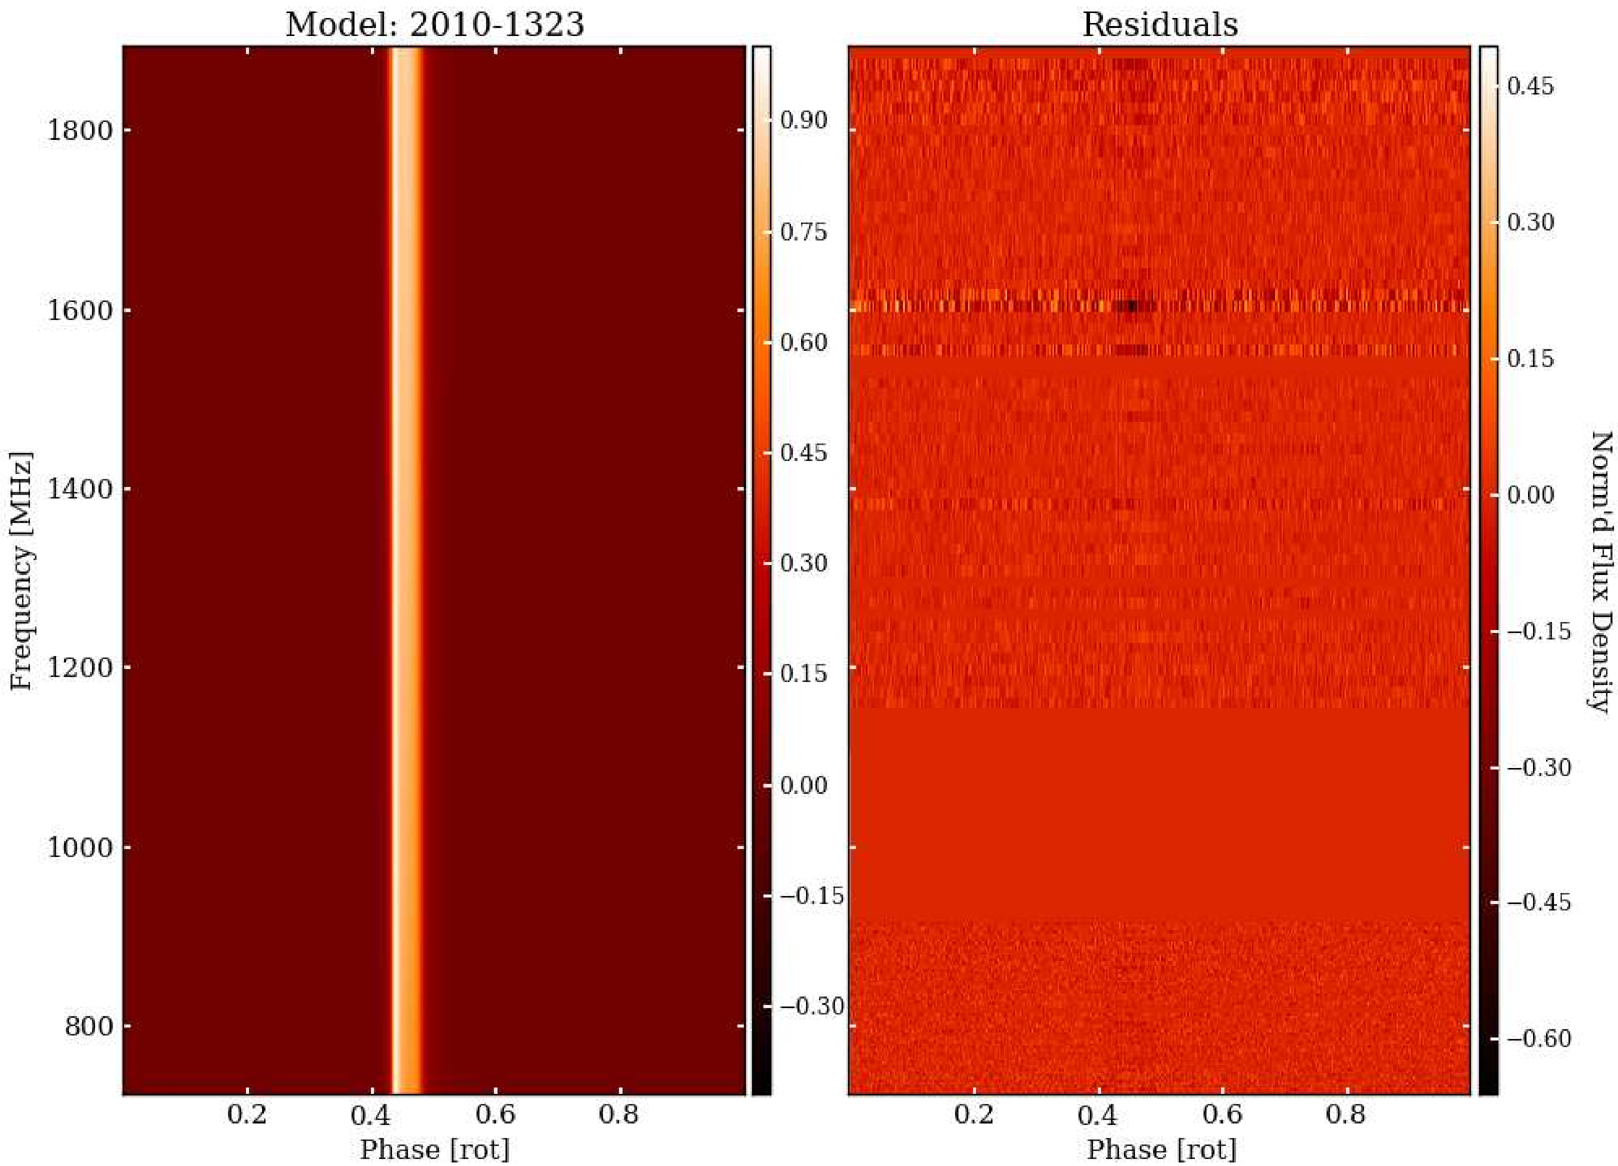

Fig. 4.217 – Fitted Gaussian component model and residuals for J2010−1323.

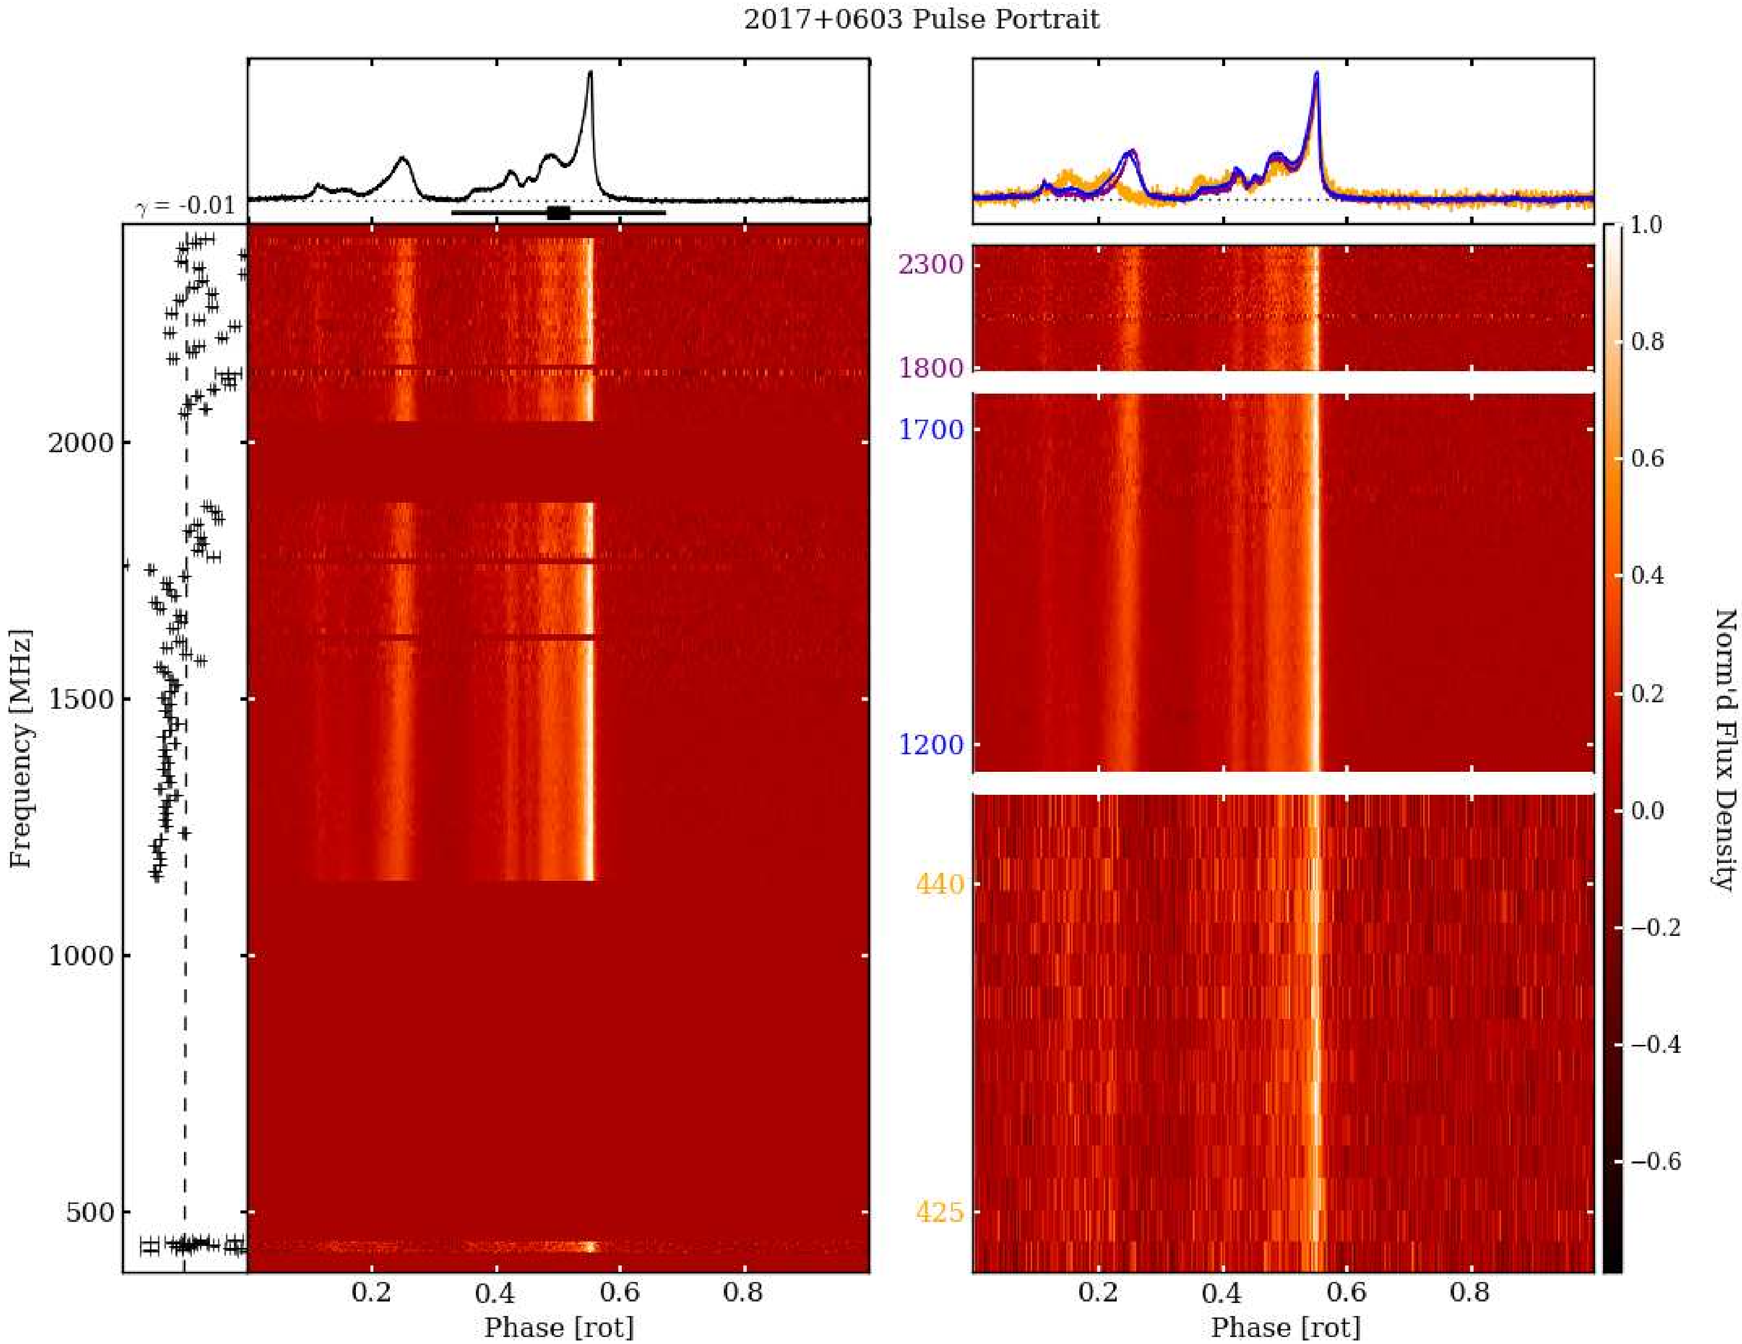

Fig. 4.218 – Concatenated portrait data and average profiles for J2017+0603.

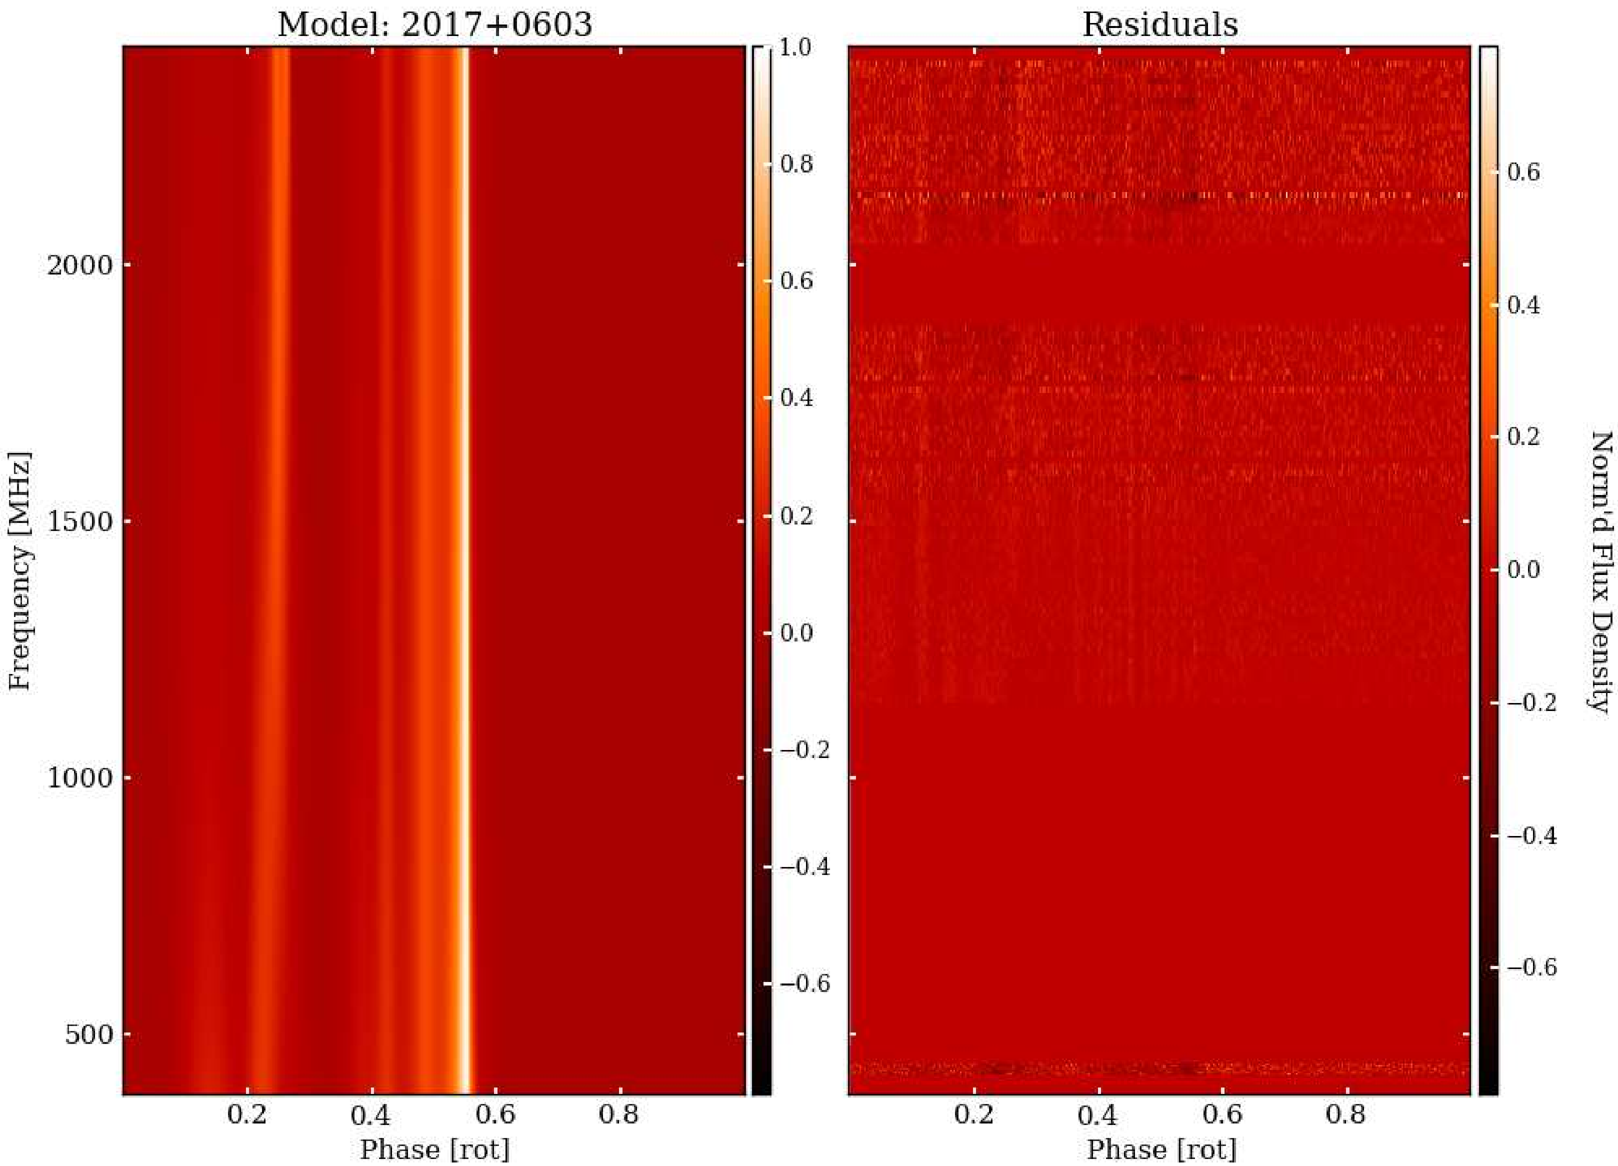

Fig. 4.219 – Fitted Gaussian component model and residuals for J2017+0603.

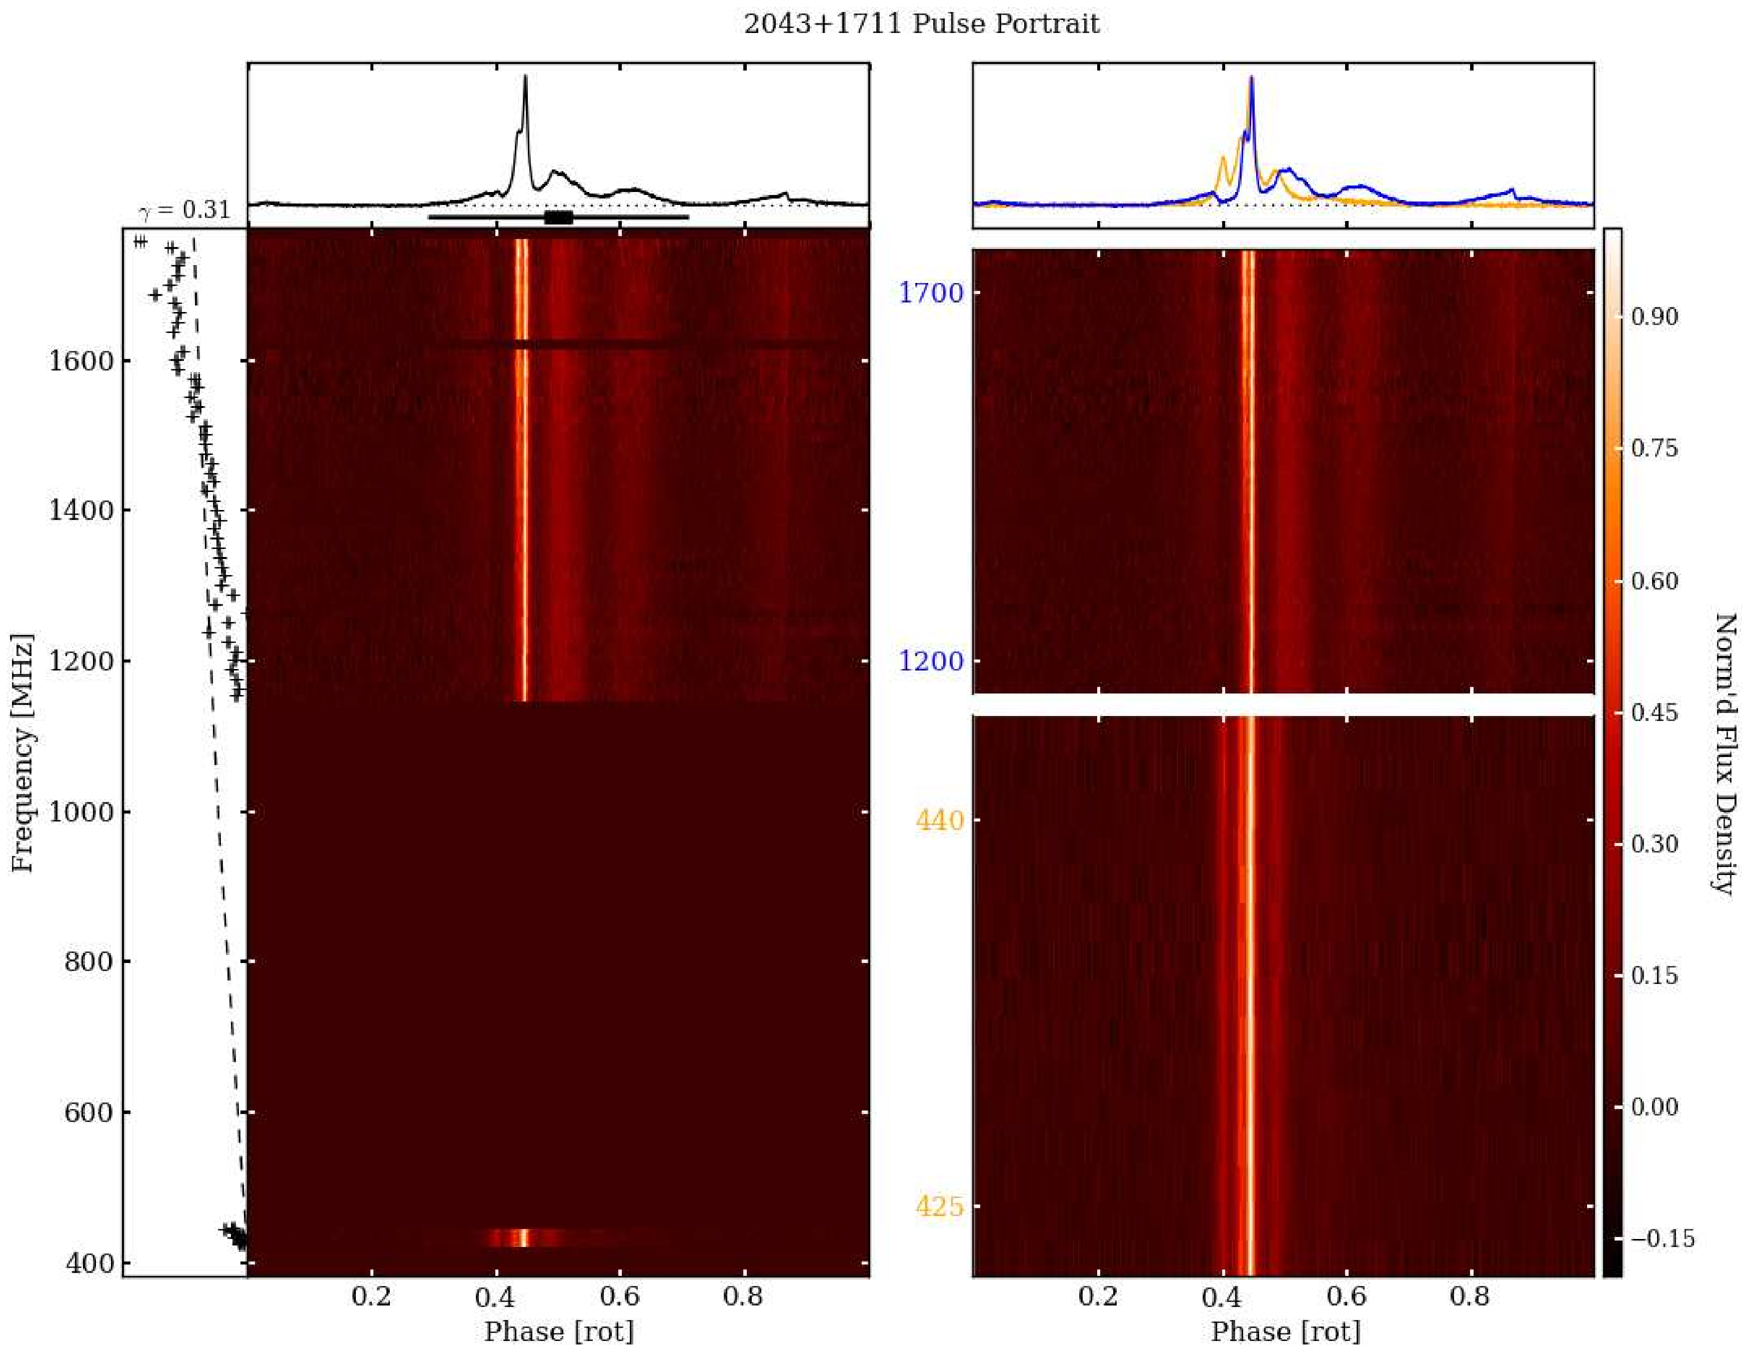

Fig. 4.220 – Concatenated portrait data and average profiles for J2043+1711.

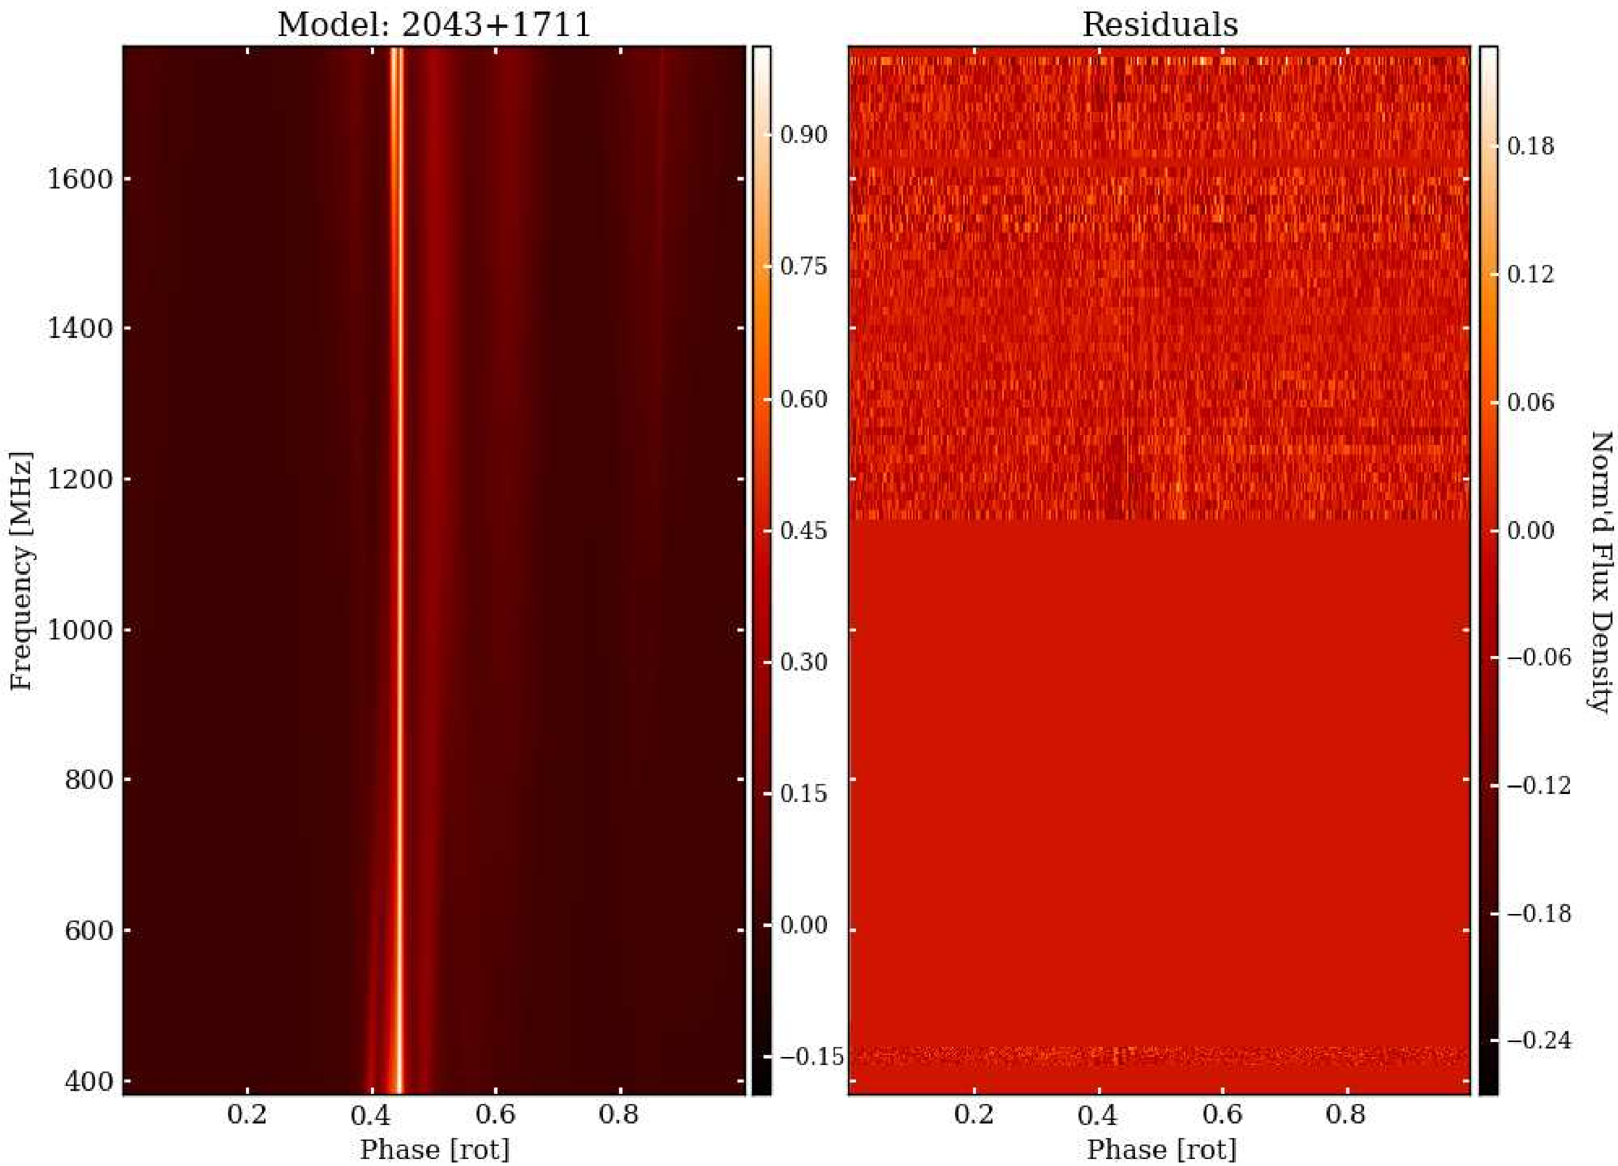

Fig. 4.221 – Fitted Gaussian component model and residuals for J2043+1711.

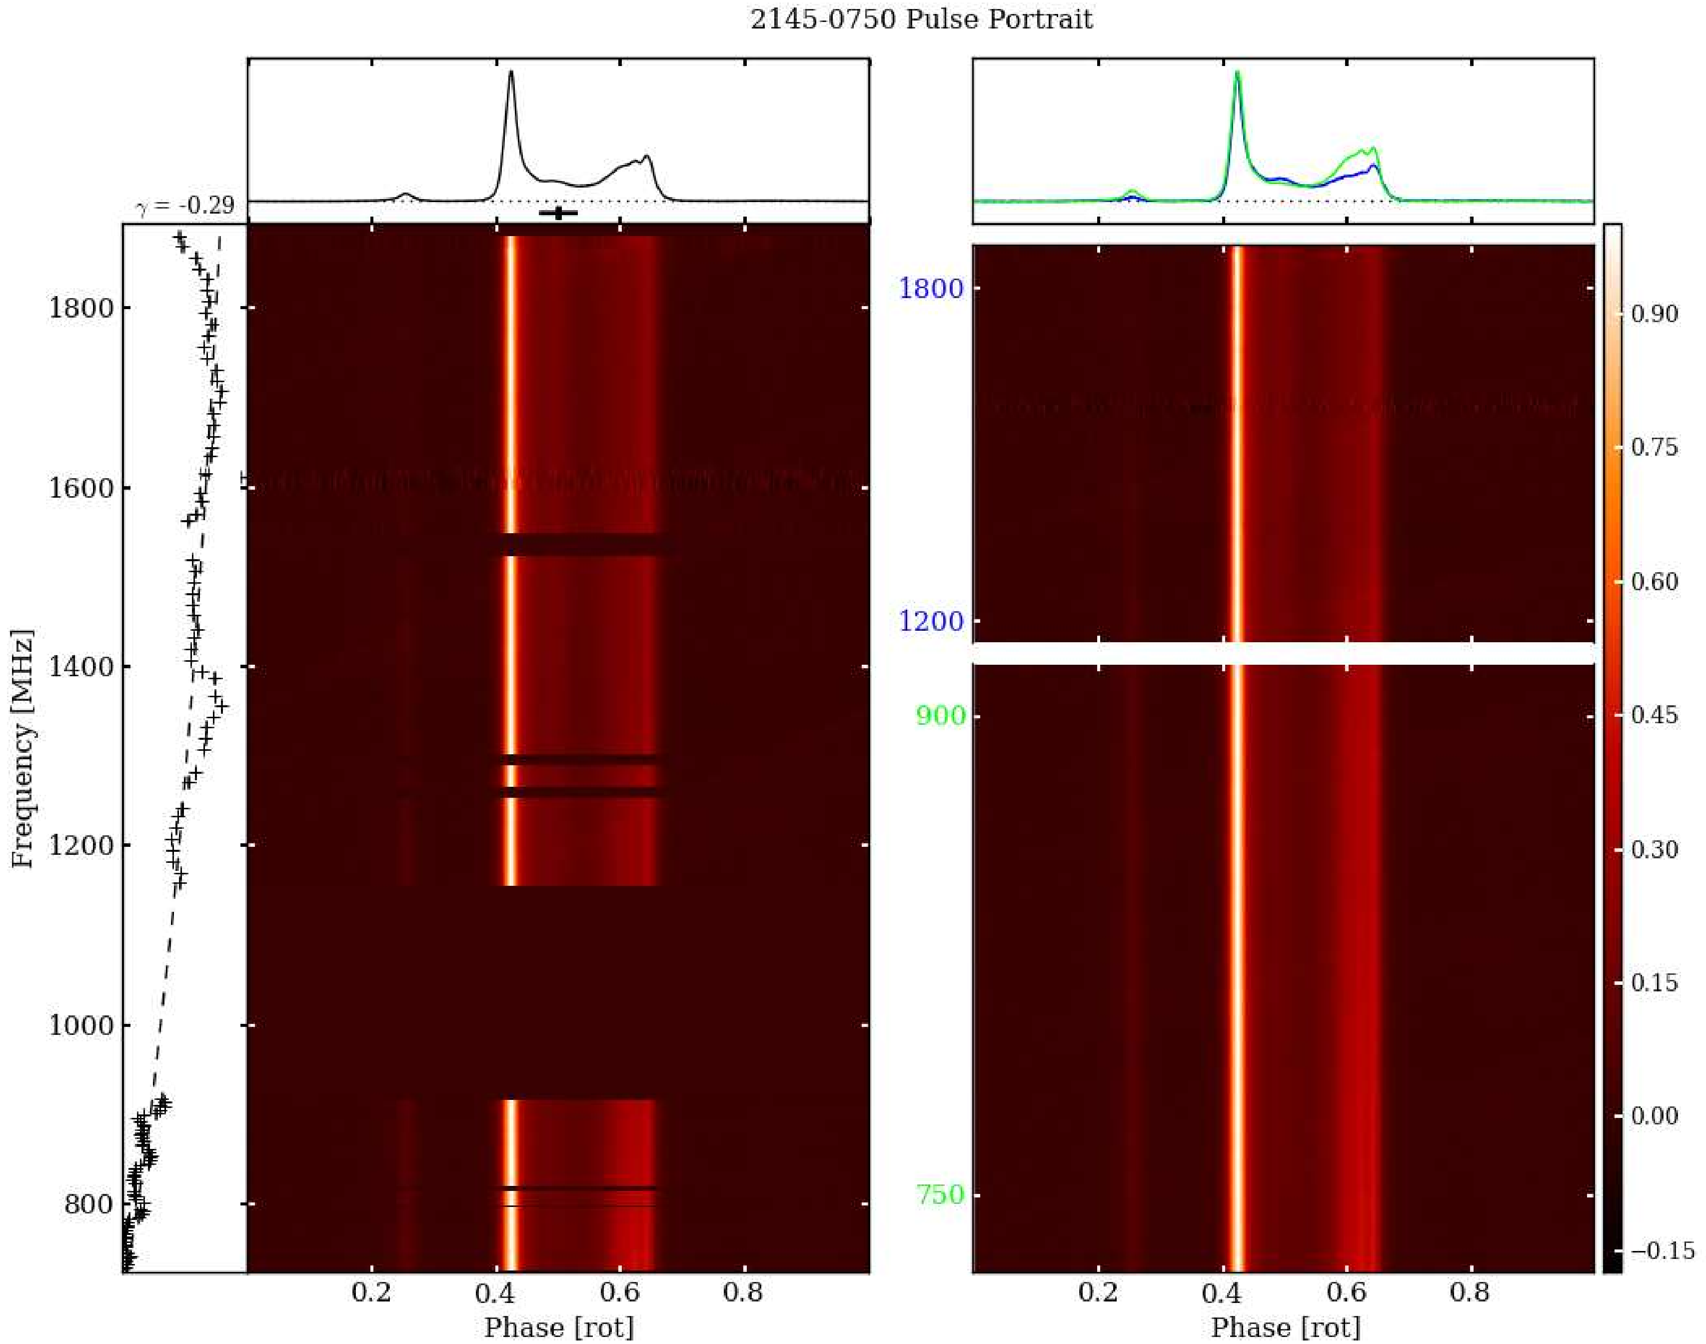

Fig. 4.222 – Concatenated portrait data and average profiles for J2145−0750.

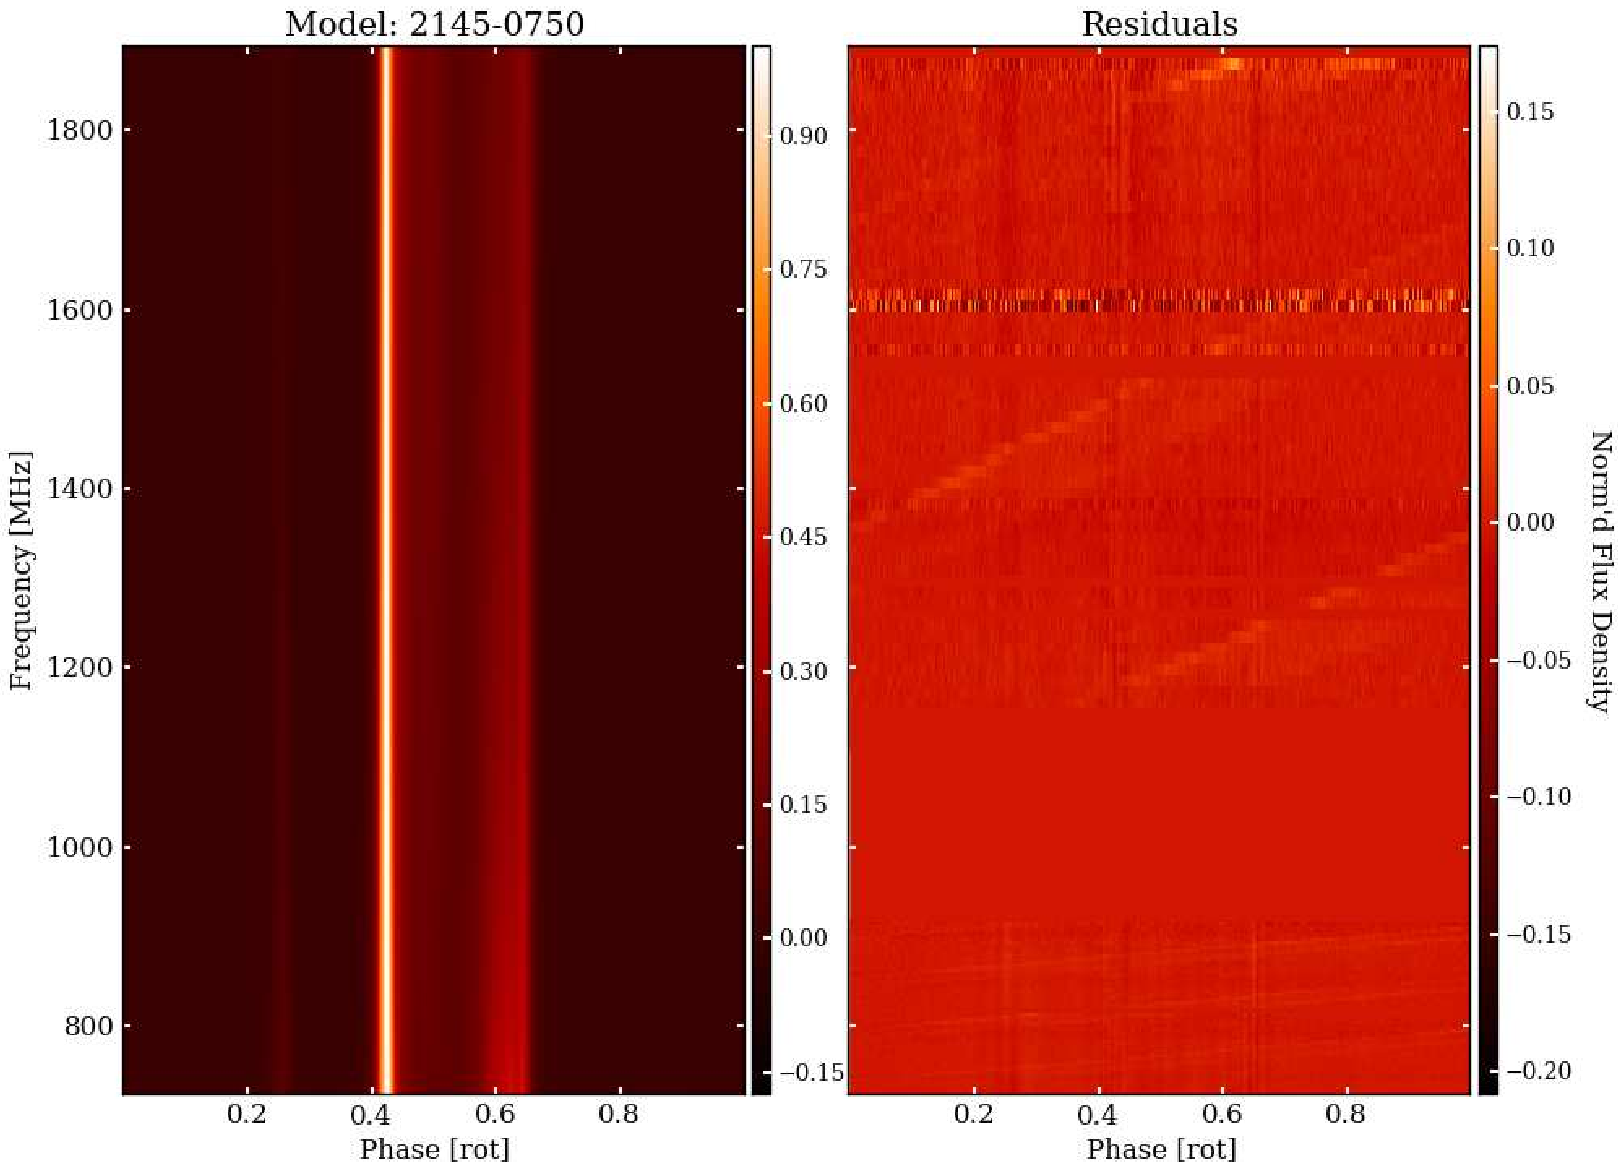

Fig. 4.223 – Fitted Gaussian component model and residuals for J2145−0750.

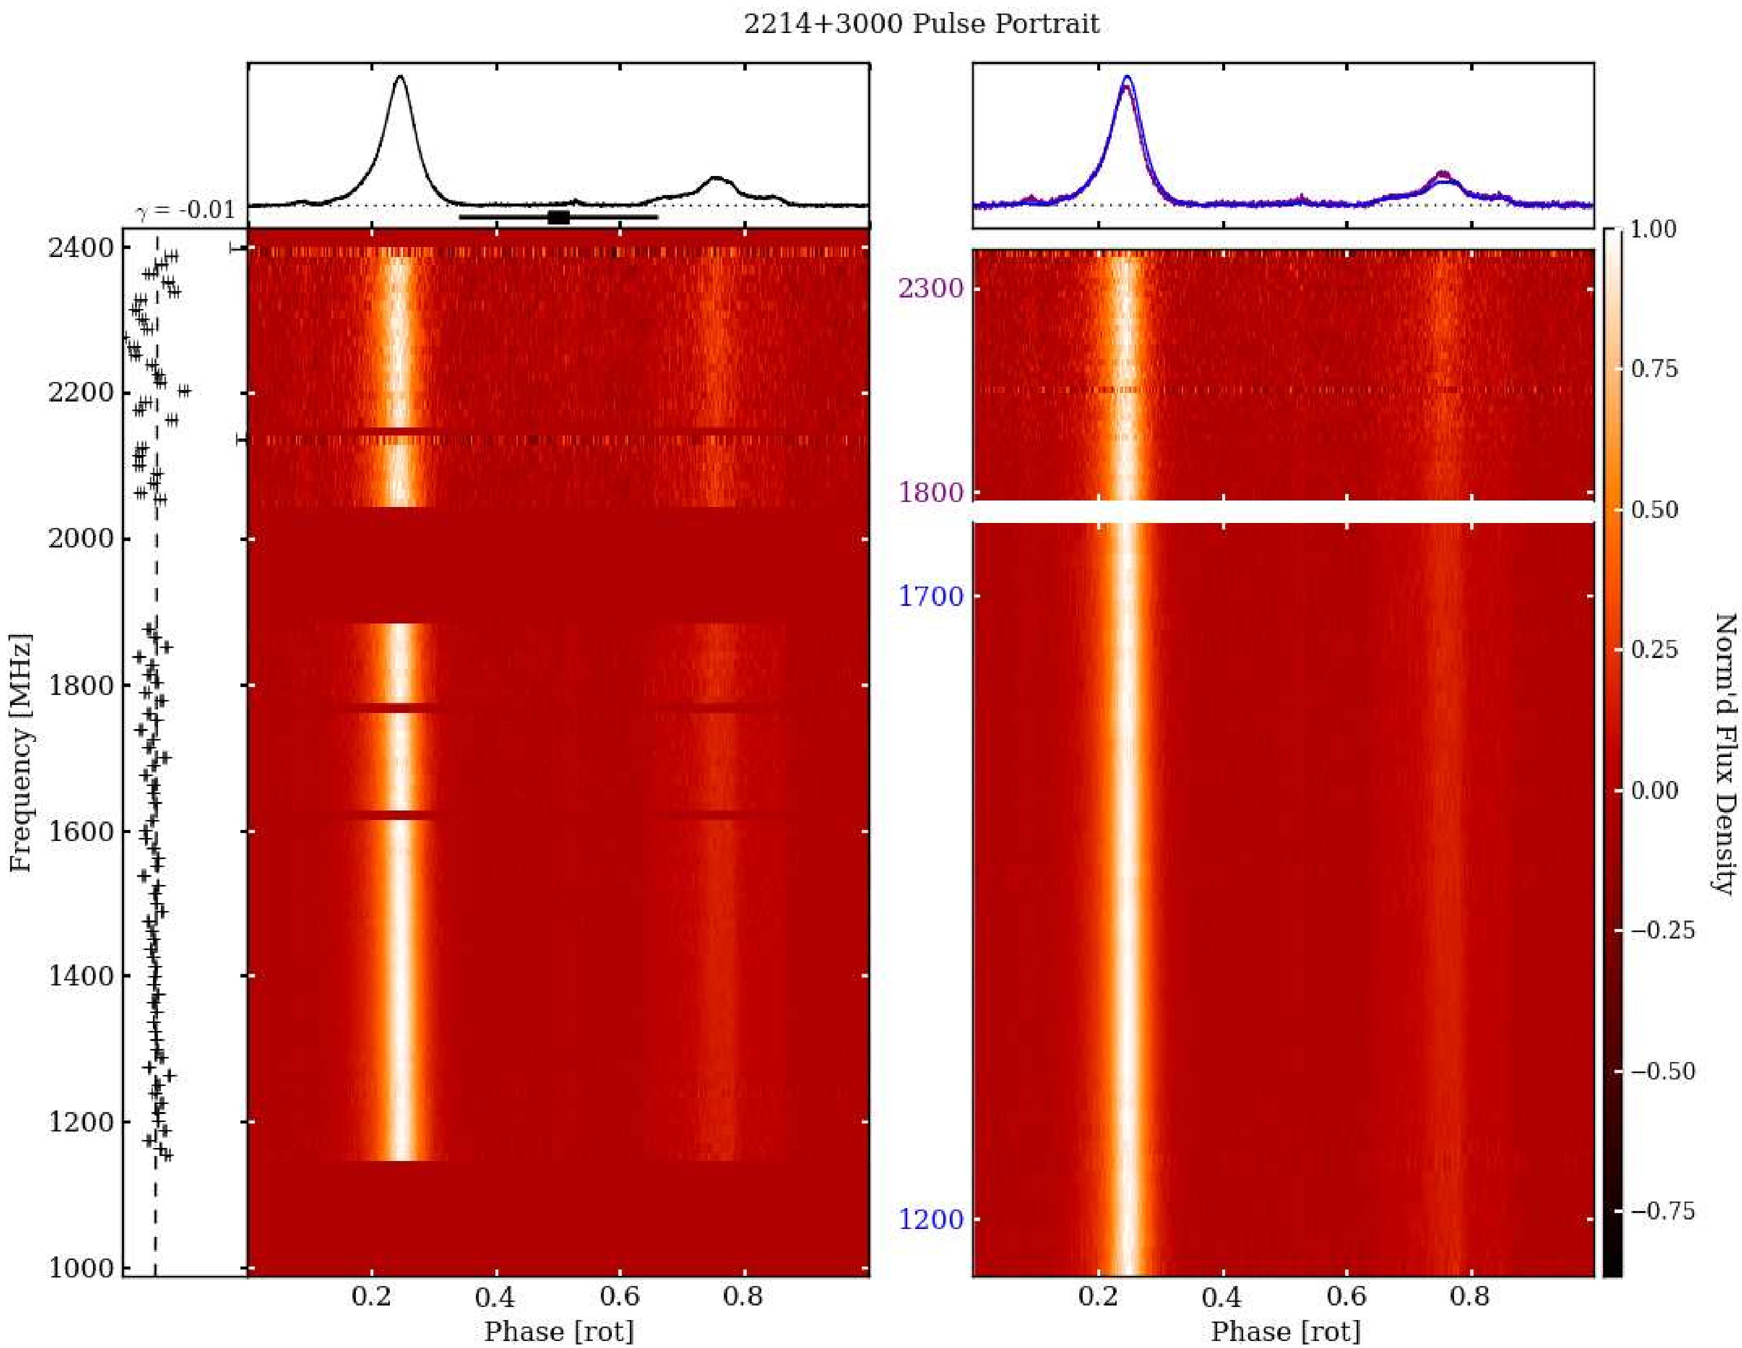

Fig. 4.224 – Concatenated portrait data and average profiles for J2214+3000.

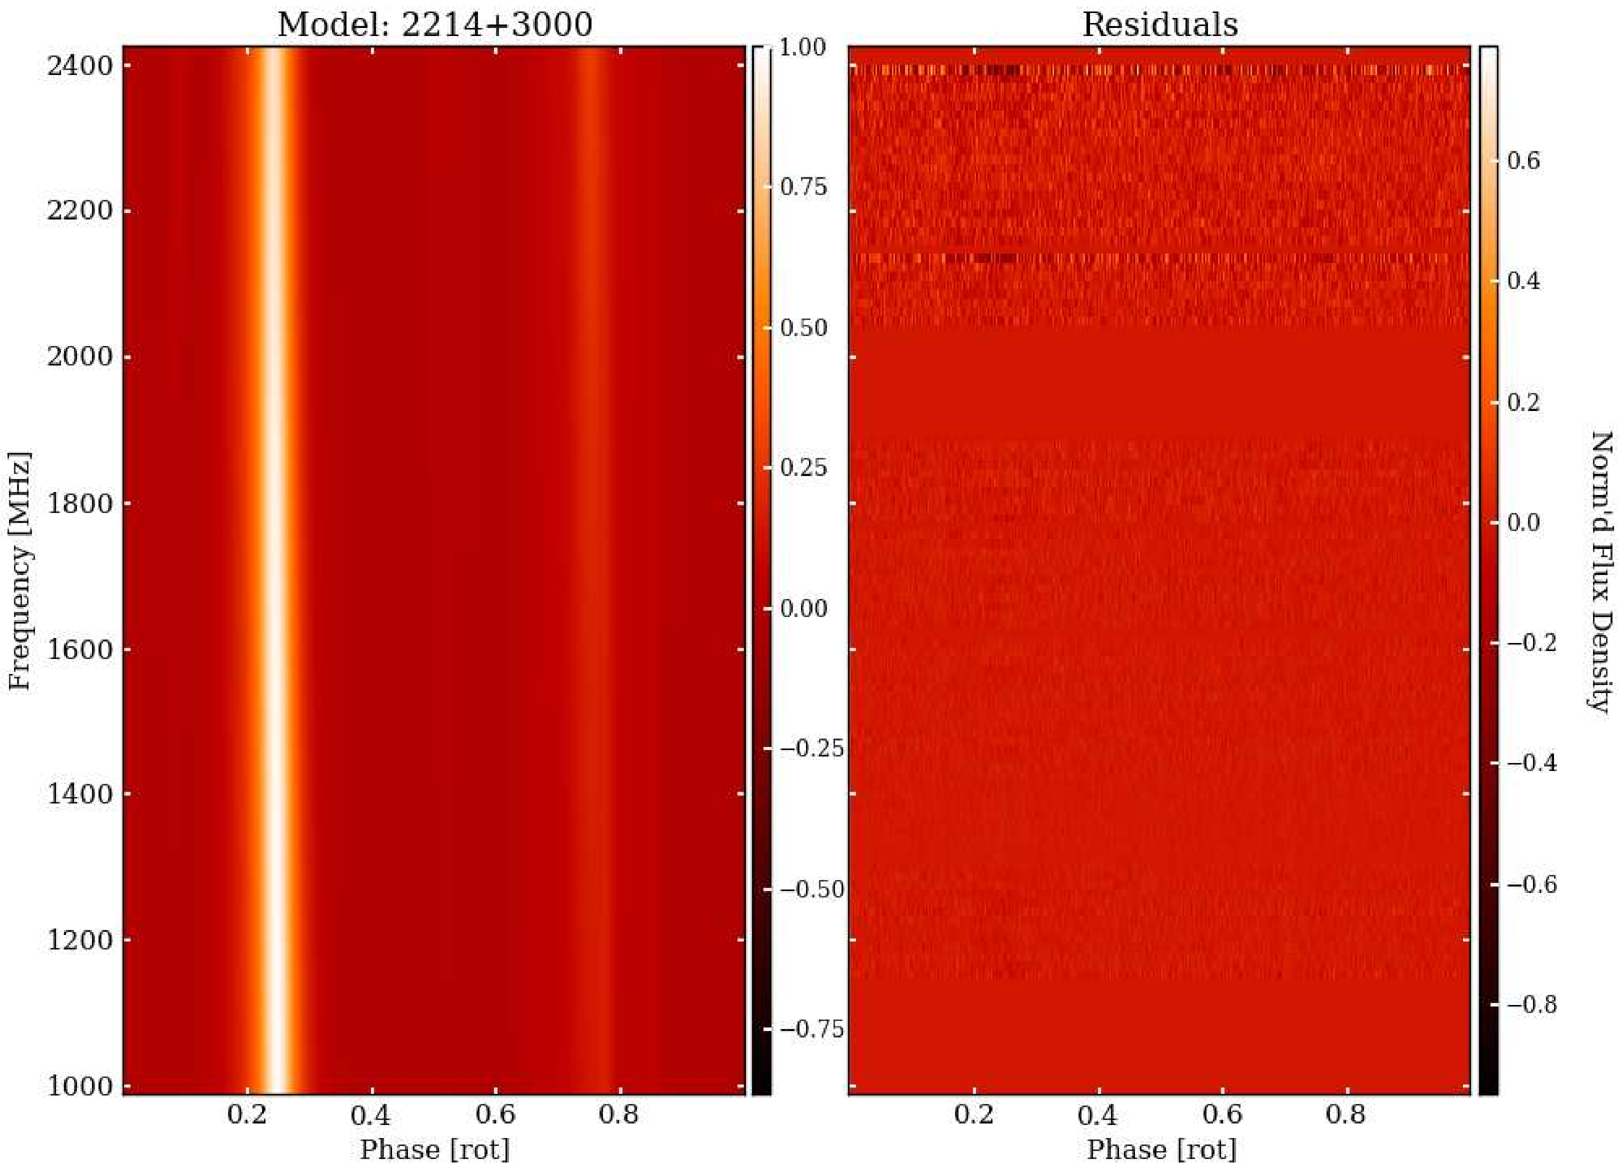

Fig. 4.225 – Fitted Gaussian component model and residuals for J2214+3000.

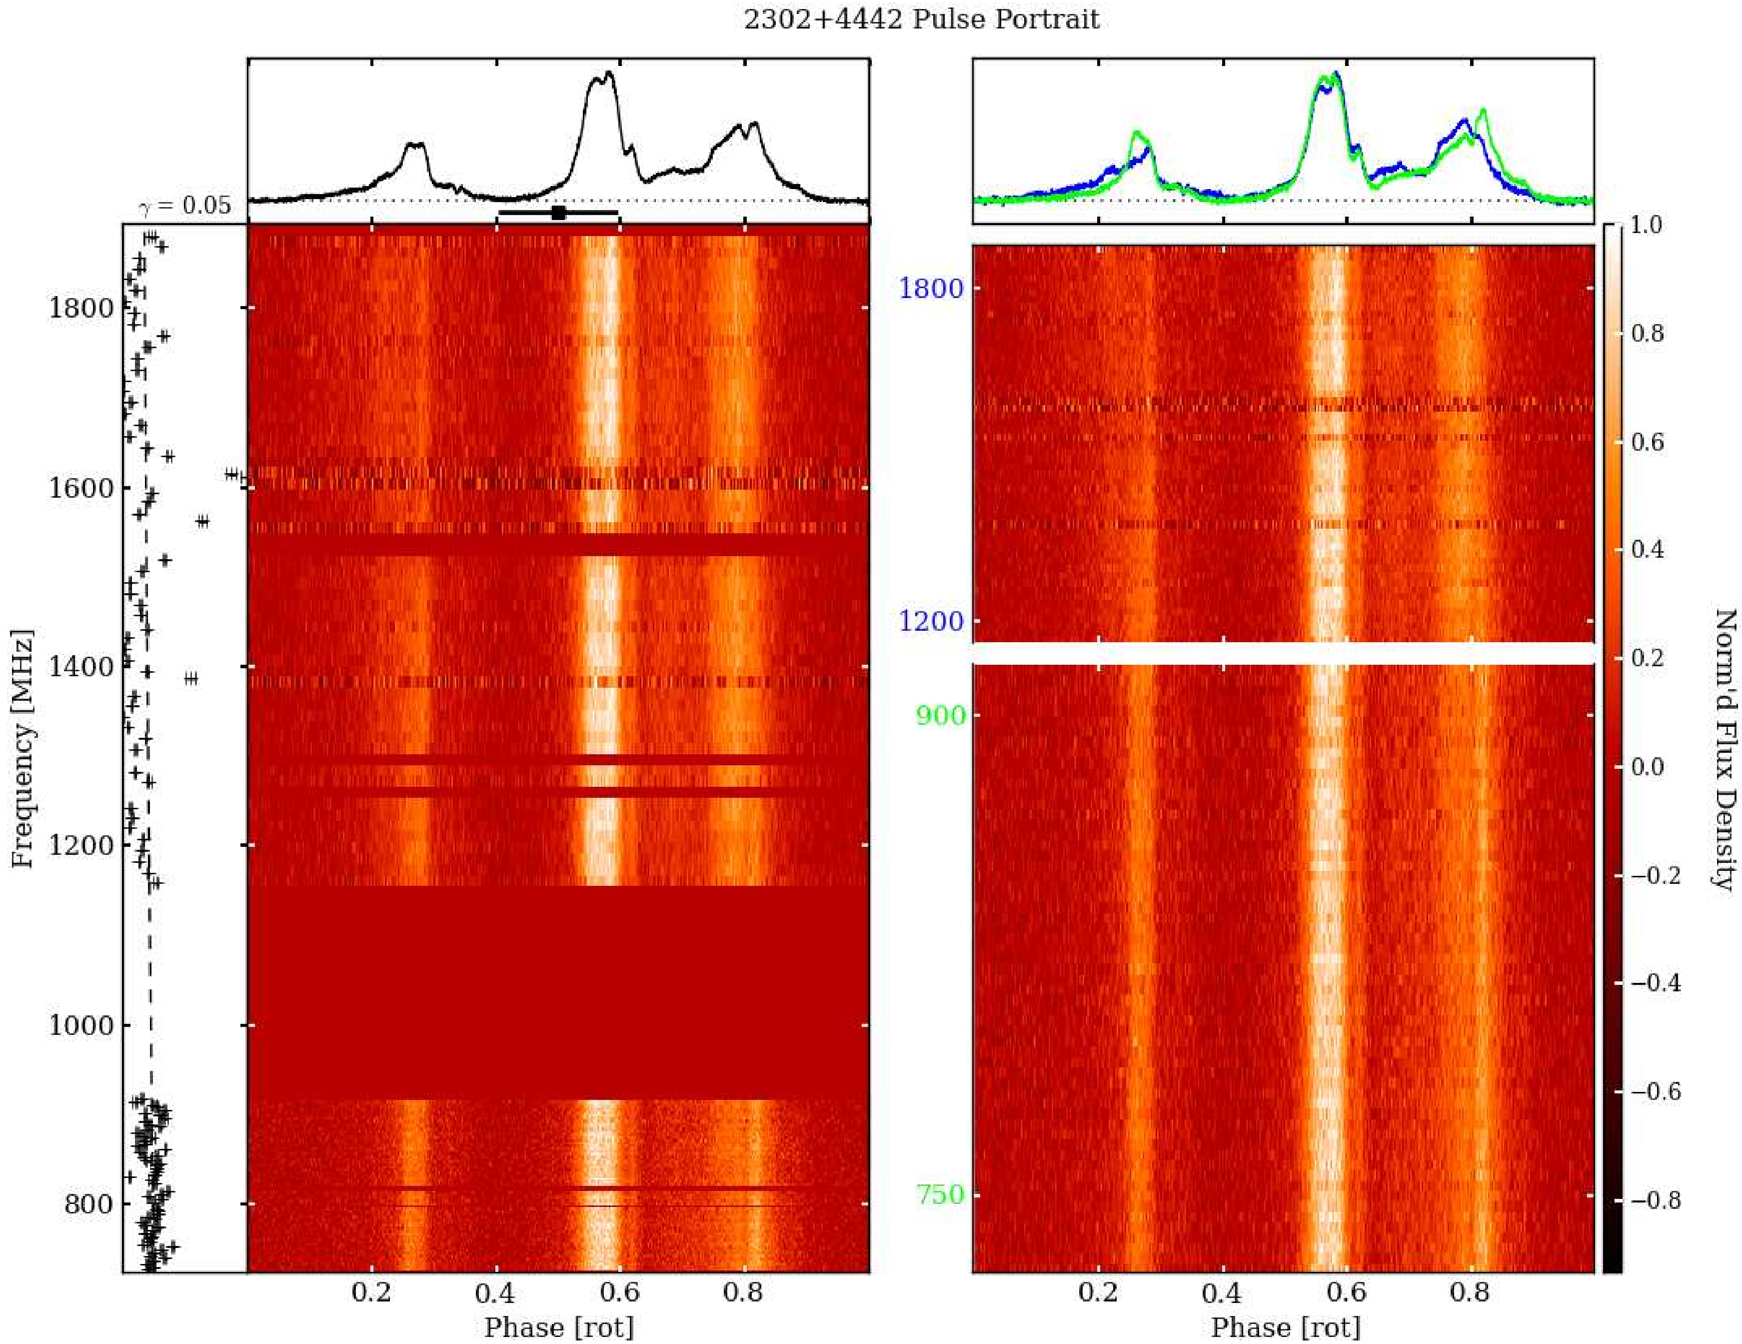

Fig. 4.226 – Concatenated portrait data and average profiles for J2302+4442.

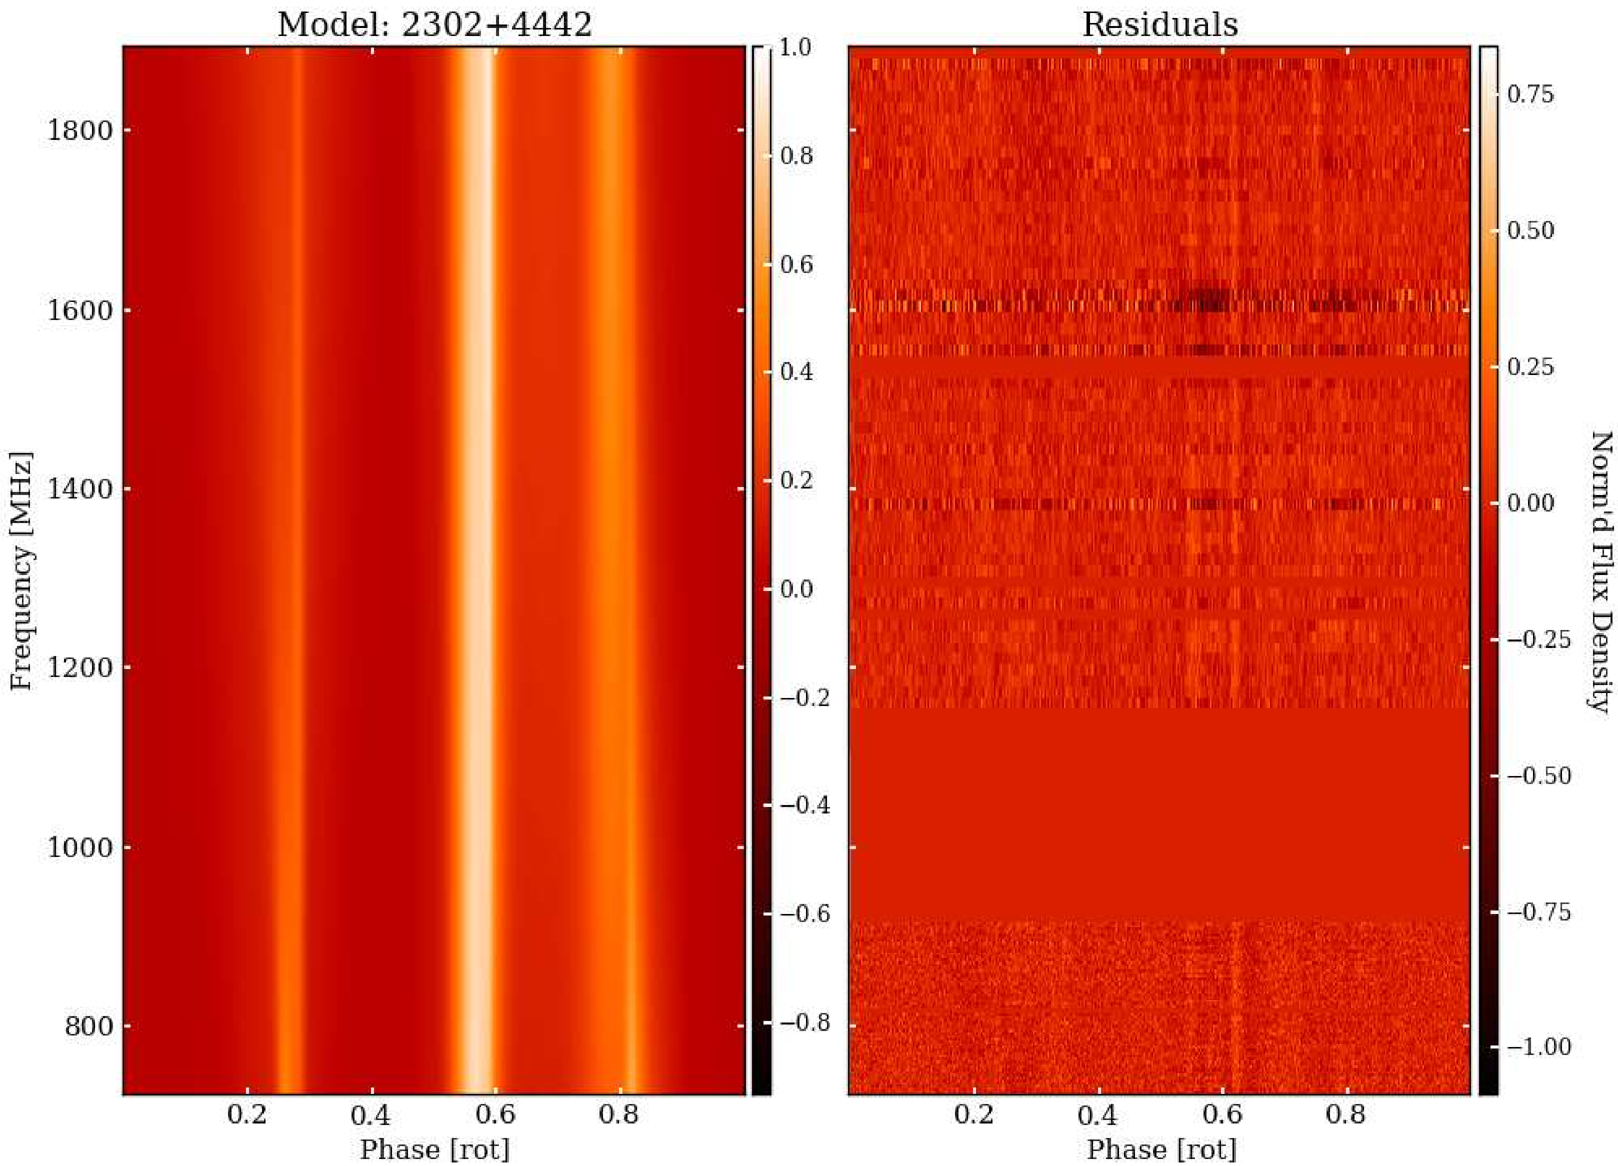

Fig. 4.227 – Fitted Gaussian component model and residuals for J2302+4442.

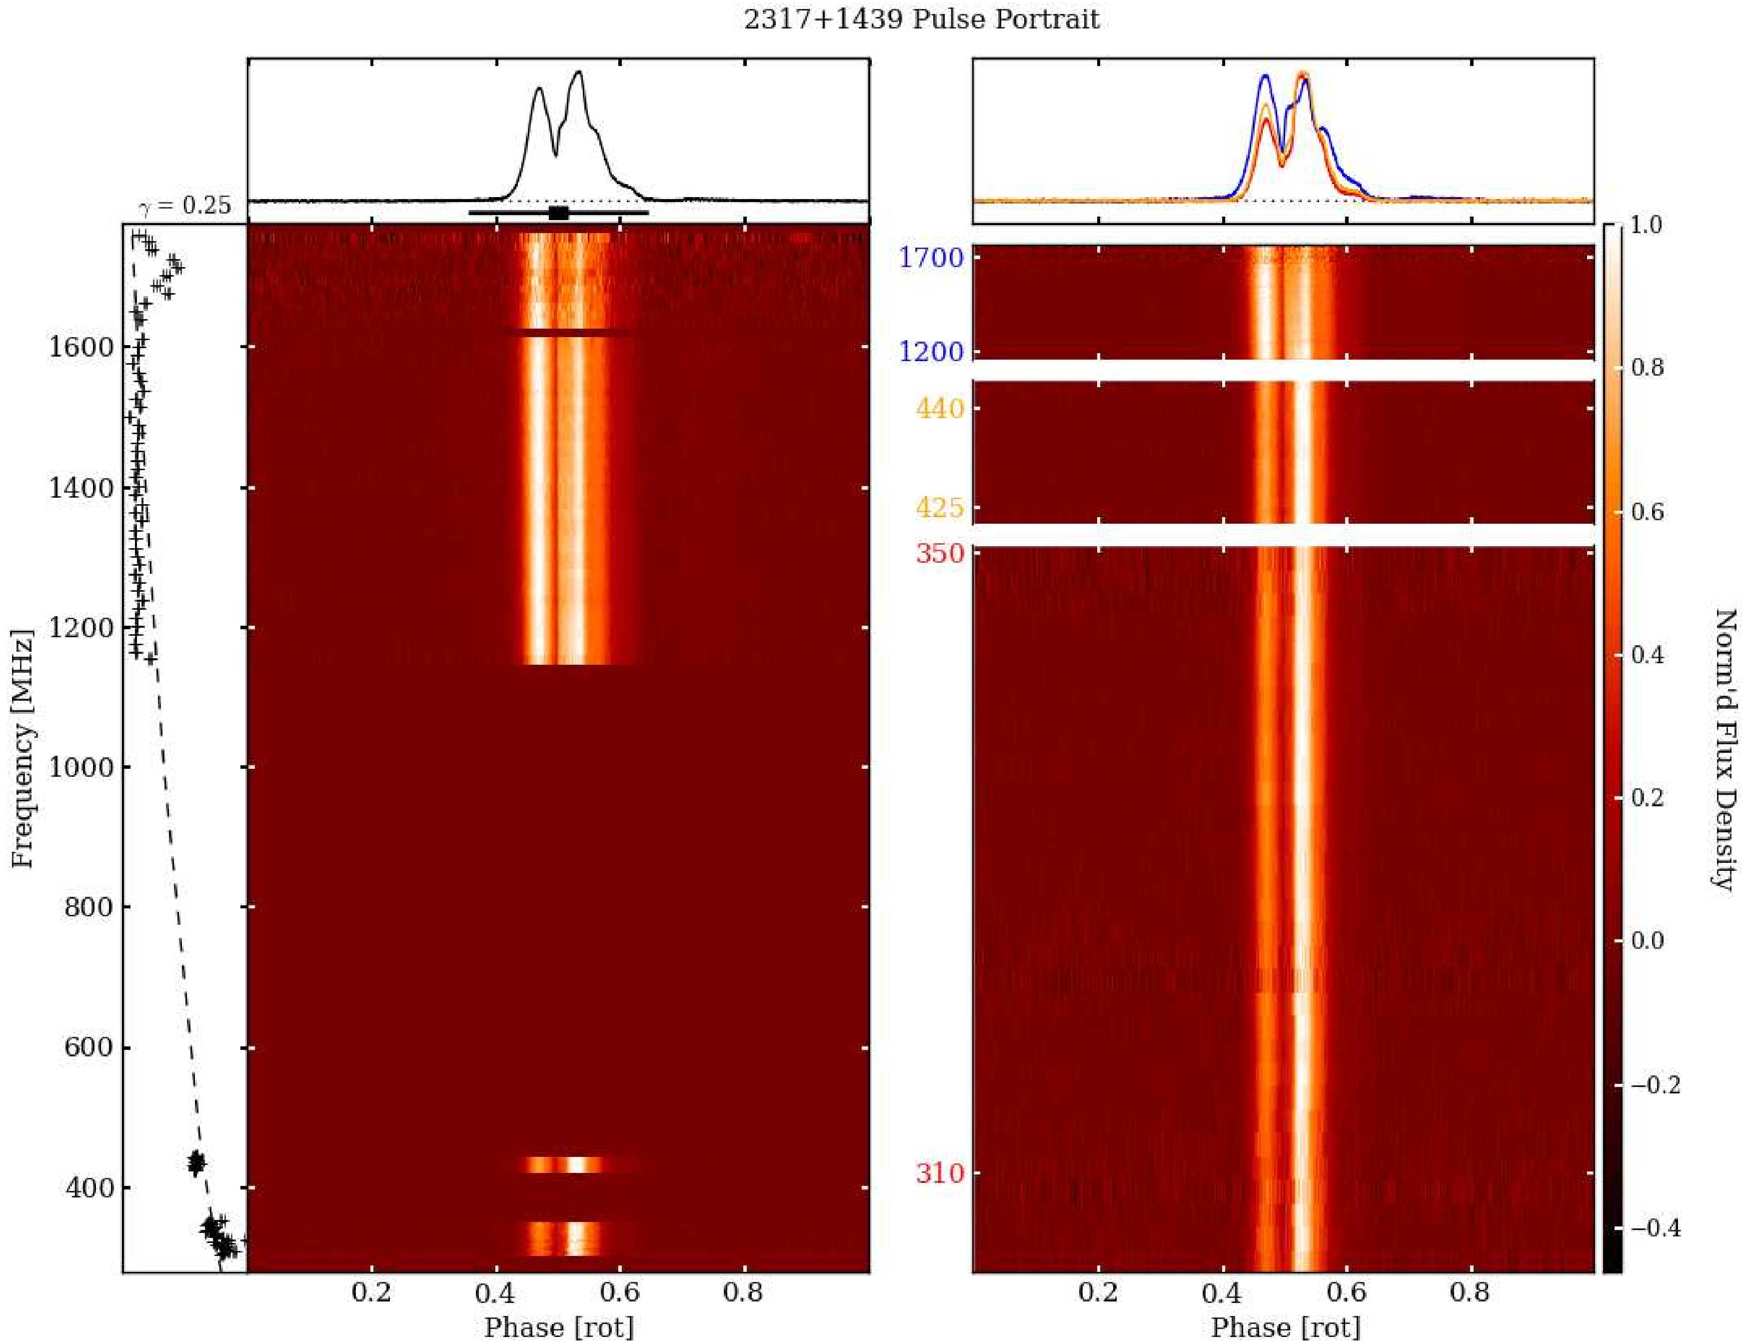

Fig. 4.228 – Concatenated portrait data and average profiles for J2317+1439.

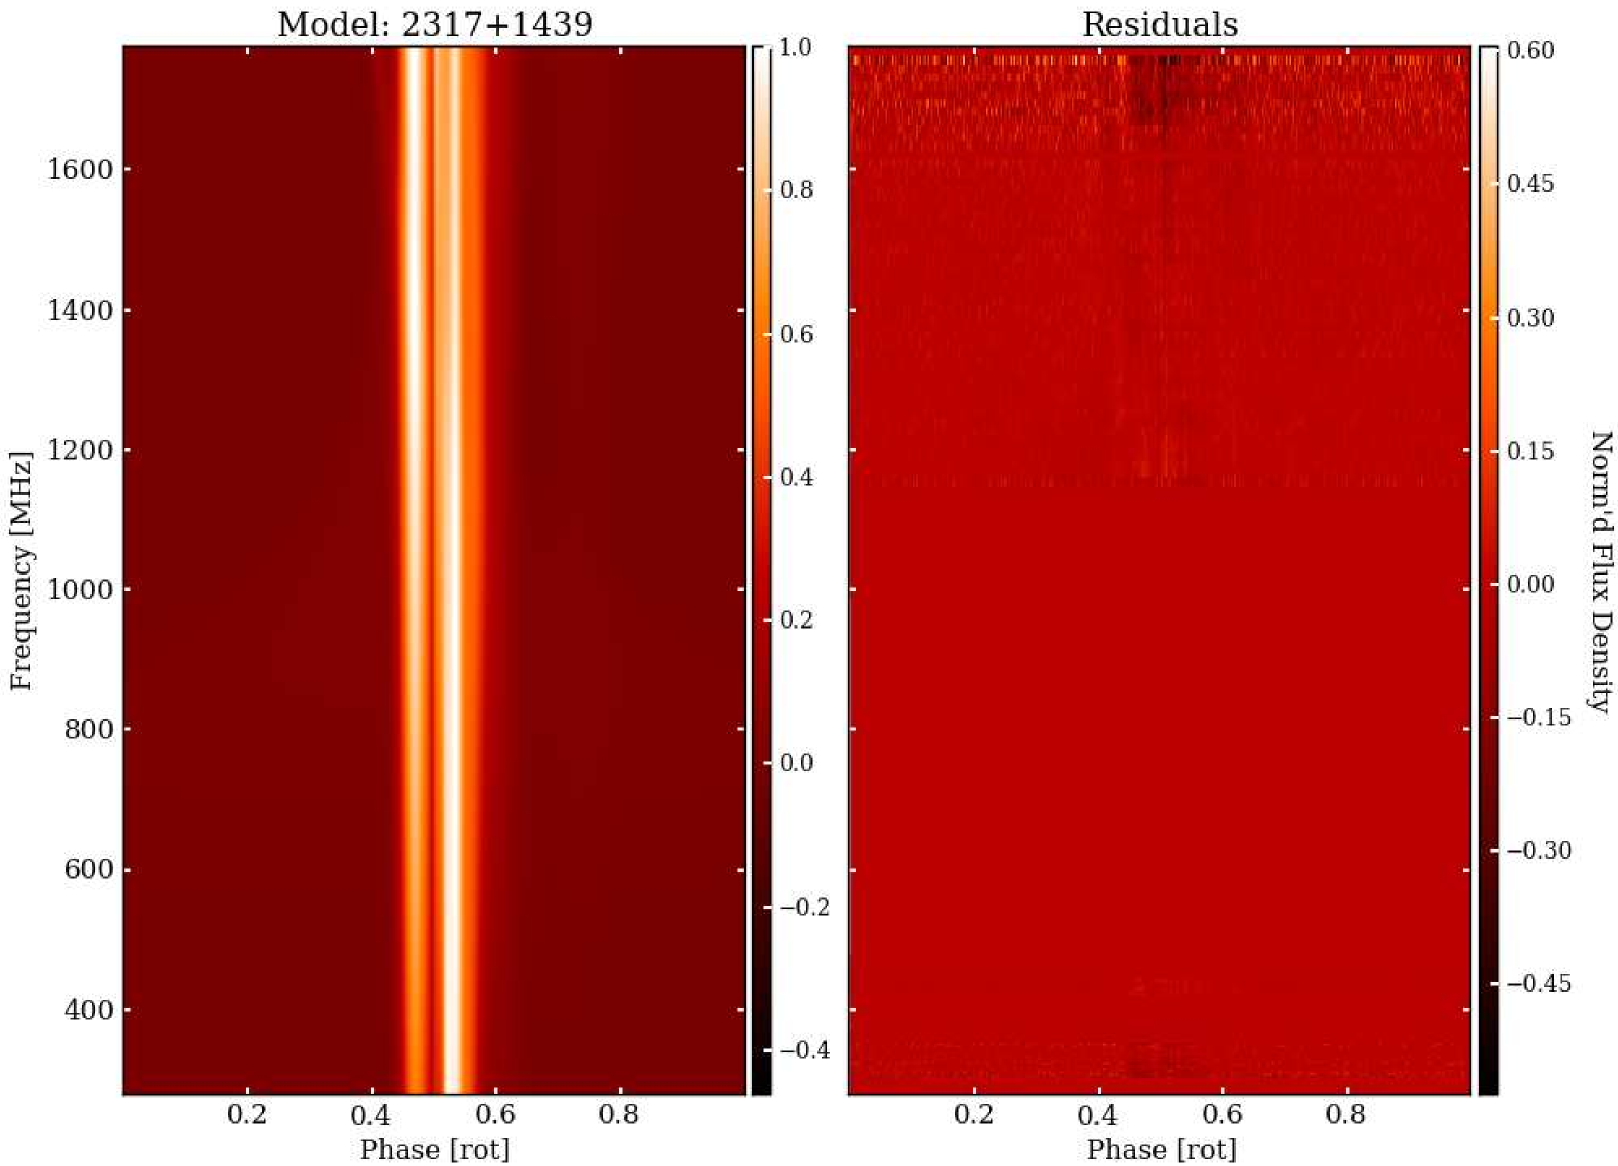

Fig. 4.229 – Fitted Gaussian component model and residuals for J2317+1439.

Chapter 5

Shapiro Delay in NANOGrav MSPs

Note: The results from this chapter will be published in tandem with other binary MSP results from the NANOGrav 9-year dataset, e.g. Fonseca et al. (in prep.).

Abstract

In this chapter, we briefly describe current, on-going efforts in the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) to extract binary millisecond pulsar (MSP) parameters — in particular, those that characterize the Shapiro delay. The Shapiro delay is a relativistic time delay that is a function of the depth of the gravitational potential traversed and the observed geometry. As such, the Shapiro delay is maximally realized in highly-inclined binary systems that have the most massive companion stars. The exemplary Shapiro delay MSP is J1614−2230, whose heavyweight neutron star was so precisely measured not only because of the fortuitous combination of hosting a 0.5 $M_\odot$ white dwarf companion and being inclined at $> 89°$, but also because of a directed set of observations (Demorest et al. 2010). Here, we describe a set of similarly specialized observations that smartly sample the orbits of NANOGrav MSPs that have unknown or not well determined Shapiro delays. These observations are included in the recent NANOGrav data release paper (Arzoumanian et al. 2015a). We demonstrate the importance of these particularly placed TOAs with MSP J2043+1711, which has a newly measured Shapiro delay thanks to our campaign. Besides the important astrophysical quantities that come from the Shapiro delay measurements, NANOGrav requires the highest possible timing precision from its MSPs, which means that accurate determinations of the binary models are a must.

5.1 Introduction

In 1964, Irwin Shapiro posited a fourth test of General Relativity (GR)[1], which predicted the existence of the now eponymous delay (Shapiro 1964). The Shapiro delay can be thought of as the excess travel time incurred by a speed-of-light signal as a result of having to traverse a gravitational potential. It arises from the apparently different proper times within the potential and not (so much) from geometric path-length differences. In the original experiment, pulsed radio waves were bounced off of Mercury and Venus when the Sun was both close and distant to the Earth–planet line of sight, and the difference between the round-trip times was compared to the theoretical predictions. Millisecond pulsar (MSP) binaries are natural extrasolar places to observe this same effect because (1) pulsars produce regular pulsed radio emission, (2) almost all pulsars that are in binaries are MSPs, (3) radio pulses are most accurately timed from the MSP population of pulsars, (4) the gravitational potential of MSP companions is sufficient to produce an observable delay, and (5) the assumption of randomly oriented orbits implies a greater number of systems that are more highly inclined, which is important for the measurement of the Shapiro delay parameters, as we will see.

For non-relativistic binaries with small eccentricities — which is the case for all pulsars studied here (the largest eccentricity being $e < 2\times10^{-4}$) — the Shapiro delay $\Delta_{\rm SD}$ takes the form,

$$\Delta_{\rm SD} = -2r\ln(1 - s\sin\Phi), \tag{5.1}$$

where r and s are the "range" and "shape" parameters, respectively, and Φ is the orbital phase as measured from the ascending node (i.e., $\Phi = \pi/2$ rad = 0.25 cycle)

[1] The first three "classical" tests explained the precession of Mercury, observed the deflection of light by the Sun, and predicted gravitational redshift.

corresponds to superior conjunction). As post-Keplerian parameters in GR, r and s take on the physically meaningful values of $r = m_{\rm c} T_\odot$ and $s = \sin i$, where $m_{\rm c}$ is the companion's mass in solar masses, $T_\odot \equiv G\mathrm{M}_\odot c^{-3} = 4.925\,490\,947$ μs is the solar mass in temporal units, and $i \in [0, 90]^\circ$ is the inclination of the binary's orbital plane relative to plane of the sky. The dashed line in Figure 5.1 shows the "full" Shapiro delay signal amplitude for $i = 89^\circ$ and $m_{\rm c} = 0.2$ $\mathrm{M}_\odot$ as a function of orbital phase. From Equation 5.1 we see that the Shapiro delay signal grows linearly with the companion mass. A Fourier series expansion of Equation 5.1 with respect to Φ shows that the coefficients of the first two harmonics of the orbital period dominate in all but the most inclined cases. The number of *measurable* harmonics will be a function of the available timing precision as well as i, but in the frequent case of only two measurable harmonics, the Shapiro delay is completely covariant with the measurement of the projected semi-major axis $x_{\rm p}$ ($= a_{\rm p} \sin i$, where $a_{\rm p}$ is the semi-major axis of the pulsar's orbit) and the eccentricity e (Lange et al. 2001; Lorimer & Kramer 2005; Freire & Wex 2010). This is evinced in Figure 5.1 by the solid lines, which show the shape and amplitude of the Shapiro delay after Keplerian orbits with $i = 40, 60, 80, 85, 89^\circ$ have been subtracted. For inclinations less than $\sim 80^\circ$, the residual structure is essentially sinusoidal when measurement uncertainties are added. For randomly oriented orbits, the probability distribution of $\cos i$ is uniform and so $\mathrm{P}(i > i') = \cos i'$; for $i' = 80^\circ$, $\mathrm{P} \approx 17\%$.

The small number of MSPs binaries where the Shapiro delay can be well measured provide three important measurements: i, $m_{\rm c}$, and the mass of the neutron star $m_{\rm p}$. Following Lorimer & Kramer (2005), a rearrangement of Kepler's Third Law gives

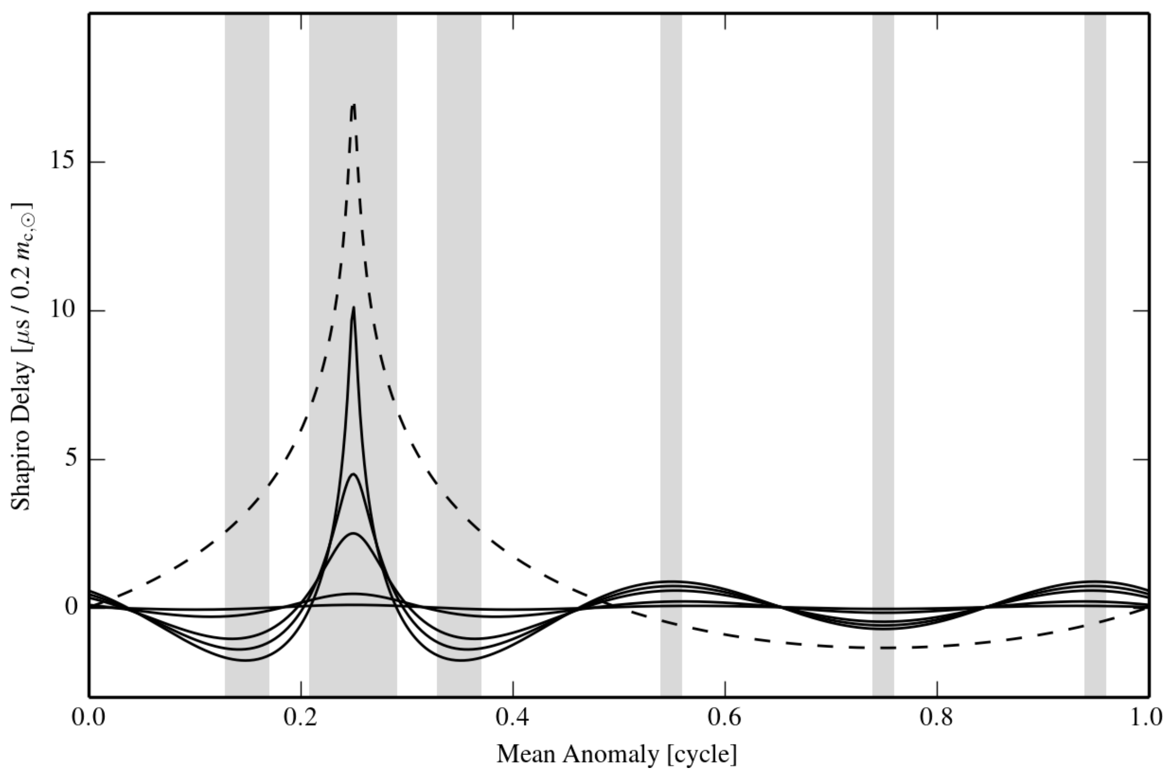

Fig. 5.1 – The shape and amplitude of the Shapiro delay signal as a function of orbital phase. The dashed line shows the "full" Shapiro delay (Equation 5.1) for $i = 89°$ and $m_{\rm c} = 0.2$ $\mathrm{M}_{\odot}$. The solid lines show the sum of the third and higher harmonics of the Shapiro delay (i.e., "post-fit", after the Keplerian orbit has been removed) also for $m_{\rm c} = 0.2$ $\mathrm{M}_{\odot}$ and $i = 40, 60, 80, 85, 89°$. The measurement of the higher harmonics is the only way to disentangle the Shapiro delay from usual Keplerian parameters. The gray regions show the orbital coverage of our observations to constrain the Shapiro delay in a subsample of NANOGrav MSPs, scaled to an orbital period of 2 days. The amplitude units are normalized by a typical carbon-oxygen white dwarf companion mass; PTA-quality MSPs usually yield better than ~ 1 μs band-averaged TOA uncertainties with typical NANOGrav observations (see Chapter 4).

the binary mass function f in solar masses,

$$f(m_{\rm p}, m_{\rm c}) = \frac{4\pi^2 x_{\rm p}^3}{P_{\rm b}^2\, T_\odot} = \frac{(m_{\rm c}\,{\rm sin}i)^3}{(m_{\rm p}+m_{\rm c})^2}, \tag{5.2}$$

where $P_{\rm b}$ is the binary's orbital period. The middle quantity is a function of two normal observables, and so $m_{\rm p}$ is determined uniquely when r and s are measurable. If not, a minimum companion mass $m_{\rm c,min}$ can be inferred by assuming a typical pulsar mass ($\sim$1.4 $M_\odot$) and setting $i = 90°$.

The values of $m_{\rm c}$ and $m_{\rm p}$ provide important constraints on the evolutionary histories of binary MSPs (e.g., Tauris et al. 2011). There are only two dozen or so well-measured neutron star masses[2], making statistical inferences about the mass distributions just now feasible (Özel et al. 2012). A collection of i measurements would allow confirmation of the long-assumed $\mathrm{P}(\cos i) \propto 1$. A bias in this distribution would arise from pulsars that more frequently beam perpendicular to or along the axis of orbital angular momentum; either scenario would warrant an interesting explanation.

Furthermore, neutron star masses provide a unique way to constrain the properties of matter at supra-nuclear densities (Lattimer & Prakash 2005; Lattimer 2012). Our earlier measurement of J1614−2230's mass (1.97 ± 0.04 $M_\odot$ at the time, which was the most-massive neutron star known) limited significantly the available parameter space for allowable neutron star matter equations-of-state (EoS; Demorest et al. 2010). Namely, many of the "softer" and more exotic EoSs were ruled out (Figure 5.2).

Aside from the direct astrophysical benefits of measuring the Shapiro delay, it is an important effect to characterize in order to minimize the timing residuals. In

[2]Jim Lattimer keeps a fairly up-to-date list of observed neutron star masses at `http://www.stellarcollapse.org/nsmasses`, as does Paulo Freire at `http://www3.mpifr-bonn.mpg.de/staff/pfreire/NS_masses.html`.

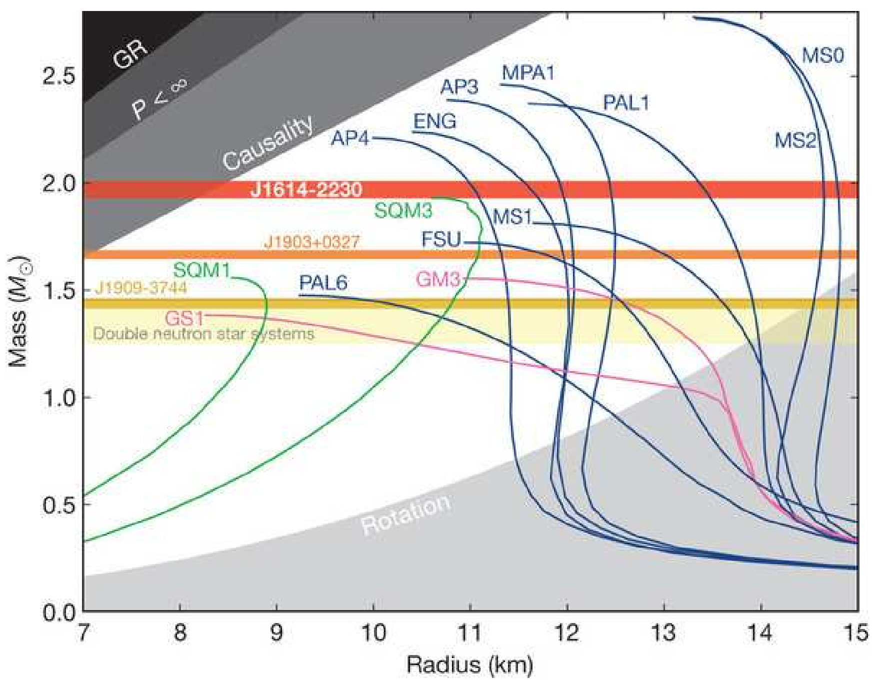

Fig. 5.2 – Neutron star equations of state. Our earlier measurement of J1614−2230's mass (red bar) ruled out a significant number of soft and exotic EoSs: blue curves represent nucleons, pink curves describe nucleons and exotic matter, and green curves are for strange quark matter. The other colored bars are mass measurements from other MSPs or double neutron stars. The gray areas are disallowed regions of the parameter space. Figure from Demorest et al. (2010).

the remainder of this chapter, we describe a set of specialized observations that we proposed and made for a subset of MSPs on behalf of the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) aimed at limiting or measuring the Shapiro delay. These observations are included in the 9-year data release that was just submitted by NANOGrav for publication (Arzoumanian et al. 2015a).

5.2 Orbital Campaigns

We proposed to supplement the ongoing, monthly timing observations in NANOGrav (§4.1.1) with a one-off set of observations to intelligently sample the orbital phases of several binary MSPs. The MSPs that were included in our subsample met any one of these criteria: (1) were newly discovered or newly added to NANOGrav, (2) had "poorly constrained" Shapiro delay parameters, and/or (3) were older MSPs that warranted a new look due to the advent of GUPPI at the Green Bank Telescope (GBT) and PUPPI at Arecibo (AO) (§4.1.1). Table 5.1 presents our list of Shapiro delay campaign MSPs; we observed these pulsars with the same telescope and observing configuration as in the normal NANOGrav timing observations.

We were granted ∼140 hr of observing time between the GBT and AO for these 12 pulsars. We scheduled our observations specifically to sample the orbit at the peaks and troughs of the post-fit Shapiro delay signal curve for each pulsar; these are the gray regions in Figure 5.1, shown scaled to an orbital period of 2 days. At the GBT, we were only granted time on each pulsar to observe conjunction and the two large troughs on either side of it, which are the most important orbital phases. All of the observations were carried out between June 2012 and January 2014.

To get a sense of what is expected, Figure 5.3 shows the post-fit Shapiro delay peak-to-trough amplitude as a function of i and m_c. The WRMS timing levels for the

Table 5.1. Shapiro Delay Campaign Pulsars

Source	P [ms]	P_b [d]	WRMS [μs]	Telescope	Notes/Reference
J0613−0200	3.06	1.2	0.229	GB	Previously limited [1]
J1012+5307	5.26	0.6	0.254	GB	Companion known; previously limited [2,3,4]
J1455−3330	7.99	76.2	1.052	GB	No measurement [5]
J1600−3053	3.60	14.3	0.203	GB	Weak detection [6]
J1643−1224	4.62	147.0	0.095	GB	No measurement [5]
B1855+09	5.36	12.3	0.414	AO	Known measurement [7,8]
J1918−0642	7.65	10.9	0.318	GB	No measurement [5]
J2017+0603	2.90	2.2	0.057	AO	New pulsar; unknown [9]
J2043+1711	2.38	1.5	0.132	AO	New pulsar; weak measurement [10]
J2145−0750	16.05	6.8	0.364	GB	Previous non-detection/limit [11]
J2302+4442	5.19	125.9	0.769	GB	New pulsar; unknown; large $m_{\rm c,min} \sim 0.3$ $M_\odot$ [9]
J2317+1439	3.45	2.5	0.259	AO	No measurement [5]

References. — [1] Hotan et al. (2006) [2] Callanan et al. (1998) [3] Lange et al. (2001) [4] Lazaridis et al. (2009) [5] Demorest et al. (2013) [6] Verbiest et al. (2009) [7] Ryba & Taylor (1991) [8] Kaspi et al. (1994) [9] Cognard et al. (2011) [10] Guillemot et al. (2012) [11] Löhmer et al. (2004a)

Note. — Demorest et al. (2013) is given as a reference for pulsars with no detected Shapiro delay if they were included in the NANOGrav 5-year data release. "Measurement" here means both Shapiro delay parameters are determined from timing observations. GB = Green Bank; AO = Arecibo. WRMS values are the (whitened) wideband residual WRMS from Table 4.5; see important notes about the timing for these pulsars in §4.2.3.

best and worst pulsars in the subsample from Table 5.1 are shown as horizontal, red, dotted lines. For randomly oriented orbits, there is a 50% *a priori* probability for a pulsar to have i in the gray region ($> 60°$). Based on this figure, we might expect that the timing of many of the pulsars in the sample would be at least somewhat informative; at the time of the observations, we predicted 1–3 new Shapiro delay measurements.

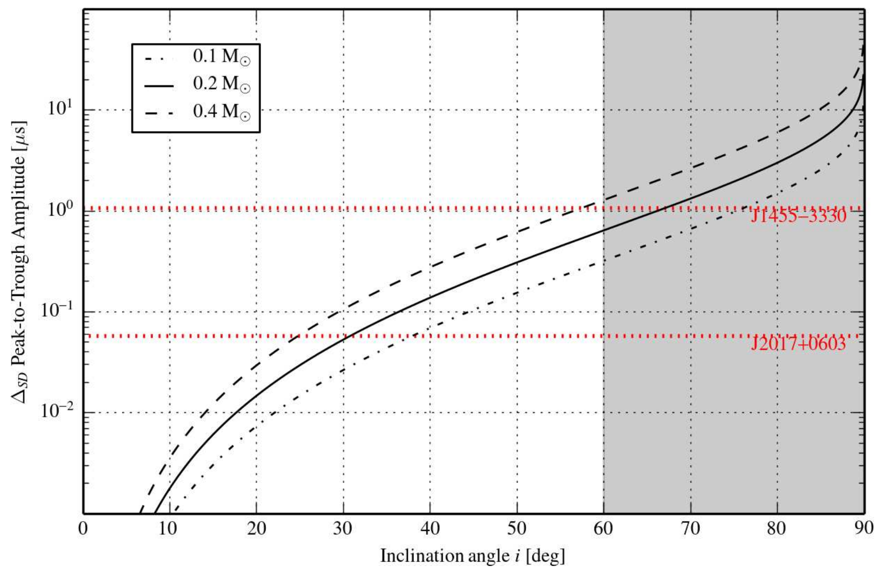

Fig. 5.3 – Post-fit Shapiro delay amplitude from Figure 5.1. The horizontal, red dotted lines are the WRMS timing residuals for the best and worst pulsars in our subsample. For the basic assumption $\mathrm{P}(\cos i) \propto 1$, a pulsar has an *a priori* probability of 50% to lie in the gray region.

5.3 Observation Results

As a follow-up to the release of the 9-year NANOGrav timing results (Arzoumanian et al. 2015a), we are carefully investigating the binary parameters for all of NANOGrav's binary MSPs. The wideband TOA dataset (Chapter 4) will be analyzed in parallel. The preliminary wideband timing results for the pulsars discussed here are as follows. The null results from the already long-term and well-timed MSPs J1455−3330, J1643−1224, and J2317+1439 will likely stand. Although the same can be said for J1012+5307, a more detailed analysis of its results that incorporates the known companion mass may yield more stringent limits. J2145−0750 and J2302+4442 still elude detection. The timing models for J0613−0200, J1600−3053, and J2017+0603 require Shapiro delay parameters[3], but only for J1600−3053 are their values somewhat constrained (J. Ellis, private communication). The masses and inclination for B1855+09 have improved in precision since earlier measurements (D. Nice, private communication). Finally, J1918−0642 and J2043+1711 comprise new, definite measurements of *both* Shapiro delay parameters[4] (E. Fonseca, private communication).

The influence of the specialized orbital campaign observations for all of the limits/detections is not obvious, but one example is enlightening. For J2043+1711, we made reduced χ^2 ($\chi^2_{\rm red}$) maps of the $\cos i$-$m_{\rm c}$ parameter space, where a timing model was fit to wideband TOAs from various subsets of observations. In each case, the best-fit timing model was found with `tempo` (including Shapiro delay parameters), and then all parameters were fixed, except for $\sin i$ and $m_{\rm c}$. These two parameters were varied over a grid of 200×200, where $\chi^2_{\rm red}$ was evaluated for each coordinate by `tempo`. In Figure 5.4 we show the results from including all observations (black), removing the TOAs from the orbital campaign observations (red), removing just the conjunction observations' TOAs (blue), and from removing n random observations (green), where n equals the number of TOAs from the conjunction observations[5].

[3] As determined by an F-test; see §4.1.1.

[4] Note that the discovery paper reporting J2043+1711 made a significant detection of the orthometric Shapiro delay amplitude h_3, which implies a large i and the presence of higher harmonics (Freire & Wex 2010; Guillemot et al. 2012).

[5] J2043+1711 was a new NANOGrav source at the start of our campaign; the number of wideband campaign TOAs is ∼30% of the total (24/77), with half of those being TOAs obtained near superior conjunction.

The contours demarcate the locus of points that are $\Delta\chi^2_{\rm red} = 1$ from the minimum $\chi^2_{\rm red}$ point. The top and side panels show the *maximum likelihood* curves for the corresponding parameter (i.e., these are cuts through the parameter space at the maximum likelihood point marked by an $\times$). The likelihood $\mathcal{L}$ was calculated from the outputted `tempo` χ^2 value as $\mathcal{L} \sim \exp(-0.5\chi^2)$.

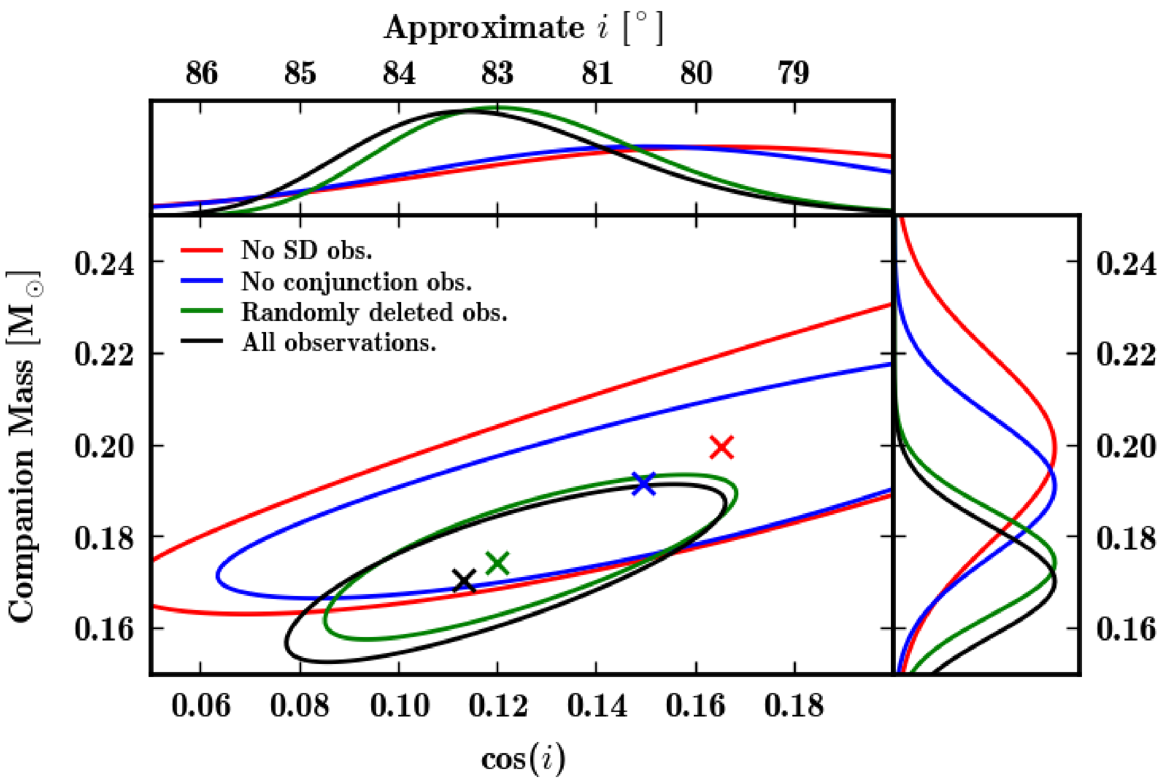

Fig. 5.4 – Reduced χ^2 maps for the Shapiro delay parameters of J2043+1711. The contours delineate a change of one in the reduced χ^2. The top and side panels plot likelihood curves for the maximum likelihood points (indicated by $\times$). The influence of the campaign observations, particularly the conjunction observations, is visible. The same number of TOAs were used to determine the blue and green curves; the only difference is that in the case of the green curves, the TOAs were randomly removed.

From the figure we can see that the TOAs obtained from the scheduled conjunction observations (all of which were within 1% of superior conjunction) play a critical role in constraining the Shapiro delay parameters, and especially the "shape" parameter. When including all observations, the 68% confidence intervals are $\cos i = 0.113^{+0.030}_{-0.024}$,

$i = 83.493^{+1.389}_{-1.742}$, and $m_c = 0.170^{+0.012}_{-0.012}$. The maximum likelihood point corresponds to a pulsar mass of the canonical value 1.35 $M_\odot$, which is in agreement with the weak constraints from Guillemot et al. (2012), but is less than their inference of a "heavy neutron star" (> 1.7 $M_\odot$) obtained when they leverage the P_b–m_c relation of Tauris & Savonije (1999) to constrain the possible values of m_c. Furthermore, the WRMS at the maximum likelihood when including all observations was 45% lower than when excluding the campaign observations. Such a drastic improvement is surely a result of the relatively highly inclined orbit ($i > 80°$), the precise timing (J2043+1711 has a median wideband TOA uncertainty of ∼200 ns), and a relatively short orbital period of just 36 hr (our carefully chosen observations covered ∼30% of the orbit, equivalent to widening the gray bars in Figure 5.1 by 50%), but the example remains instructive. An analysis of the channelized TOA data set using all observations gives essentially the same values for the maximum likelihood points and confidence intervals, but the error ellipse is somewhat more inclined (J .Ellis, private communication). It is not clear why our simple likelihood analysis of the wideband TOAs here should show less covariance in the parameters.

5.4 Project Future

The robust determination of the Shapiro delay parameters in the NANOGrav 9-year data set is still very much a work in progress between a number of people in NANOGrav. The end goal here is to amass a set of robust limits and measurements, no matter how imprecise, for all of the NANOGrav MSPs. The first step for our "wideband" version of the 9-year data set is to ensure robust noise modeling (see §4.2.3) and compare these sensitive non-linear binary parameters to those from the standard TOA analyses. One other open question is how to best place prior probability distributions on $m_{\rm c}$, $\sin i$, and $m_{\rm p}$, to get robust inferences on these parameters from a full Bayesian analysis. After these things are resolved, we will be able to quickly infer the statistical consistency of the $\cos i$ distribution with the assumed flat distribution.

As it turns out, our expectation of 1–3 new Shapiro delay detections from our MSP subsample was on point. It is worthwhile to ensure that we have a fair amount of coverage around the precise phase of conjunction, since this appears to be the best way to meaningfully constrain the Shapiro delay, and especially in low inclination systems. Measuring the parametric Shapiro delay amplitude (h_3, Freire & Wex 2010) will determine whether or not any of the remaining pulsars have Shapiro delay harmonics in their residuals, leading to possible additional campaigns. The addition of new Shapiro delay measurements is slow, since one is limited by the available timing precision and the time it takes to cover an orbit. However, we have shown that routine Shapiro delay campaigns are a straightforward way to obtain better timing and quick scientific returns for any new MSPs that get added into the array.

Chapter 6

Simultaneous Multi-band Radio & X-ray Observations of the Galactic Center Magnetar SGR 1745−2900

Note: This chapter comprises a paper of the same name that was accepted by *The Astrophysical Journal* on 8 June 2015, Pennucci et al. (2015, arXiv:1505.00836).

6.1 Project Background: The NANOGrav PIRE

The NANOGrav PIRE program[1] facilitated a travel program to "partner IPTA institutions", which I elected to participate in. After having attended two conferences in Sardinia near the site of the nascent 64-m Sardinia Radio Telescope, I decided to collaborate with Italian astronomers in the hopes of gaining insight into the first steps of a brand new telescope, as well as work on my language skills. I was enabled to spend three months at the Osservatorio Astronomico di Cagliari from January–March 2014. Unfortunately, the telescope's dual L/P-band receiver and wideband backend were not yet fully functional, so nothing materialized from the prospect of me exploring the new data with my wideband timing/modeling scheme. However, I was given this interesting magnetar project to work on, which turned out to have a host of unexpected results.

This chapter represents the end product from this PIRE exchange, and it happens to fit quite nicely within the wideband context of this dissertation. J1745−2900 is a slow pulsar (a magnetar), whereas I had only considered MSPs up until this point. All of the interesting ISM measurements in this chapter were made possible by the broadband radio observations with GUPPI. The large scattering timescale, the variable scattering index, and the potentially frequency-dependent dispersion measure all required innovation and that I augment my modeling/timing codes.

Besides the sections on the X-ray reduction/spectral analysis (completed by P. Esposito) and the discussion of the magnetar in context (written by N. Rea), the work is entirely mine. The co-authors on the paper are: Andrea Possenti, Paolo Esposito, Nanda Rea, Daryl Haggard, Frederick Baganoff, Marta Burgay, Francesco Coti Zelati, GianLuca Israel, and Toney Minter.

[1] `http://nanograv.org/pire/` (NSF PIRE Grant 0968296)

Abstract

We report on multi-frequency, wideband radio observations of the Galactic Center magnetar (SGR 1745−2900) with the Green Bank Telescope for ∼100 days immediately following its initial X-ray outburst in April 2013. We made multiple simultaneous observations at 1.5, 2.0, and 8.9 GHz, allowing us to examine the magnetar's flux evolution, radio spectrum, and interstellar medium parameters (such as the dispersion measure (DM), the scattering timescale and its index). During two epochs, we have simultaneous observations from the *Chandra* X-ray Observatory, which permitted the absolute alignment of the radio and X-ray profiles. As with the two other radio magnetars with published alignments, the radio profile lies within the broad peak of the X-ray profile, preceding the X-ray profile maximum by ∼0.2 rotations. We also find that the radio spectral index γ is significantly negative between ∼2 and 9 GHz; during the final ∼30 days of our observations $\gamma \sim -1.4$, which is typical of canonical pulsars. The radio flux has not decreased during this outburst, whereas the long-term trends in the other radio magnetars show concomitant fading of the radio and X-ray fluxes. Finally, our wideband measurements of the DMs taken in adjacent frequency bands in tandem are stochastically inconsistent with one another. Based on recent theoretical predictions, we consider the possibility that the dispersion measure is frequency-dependent. Despite having several properties in common with the other radio magnetars, such as $L_{\mathrm{X,qui}}/L_{\mathrm{rot}} \lesssim 1$, an increase in the radio flux during the X-ray flux decay has not been observed thus far in other systems.

6.2 Introduction

Magnetars are exotica among the exotic: whereas other pulsars are sustained by their stored angular momentum, the primary energy source that powers this special class of objects is likely the neutron star's immense magnetic field (Mereghetti et al. 2015). The field strengths take on the highest values ever inferred, typically $> 10^{12}$ G and even up to $\sim 10^{15}$ G. According to the McGill Online Magnetar Catalog[2] (Olausen & Kaspi 2014), there are 28 known magnetars, of which only four have displayed pulsed radio emission.

SGR 1745−2900 (J1745−2900, hereafter) is the most recent addition to the small collection of magnetars with observed pulsed radio emission (the "radio magnetars", to which we will refer by their PSR names: J1809−1943 (XTE 1810−197), J1550−5418 (1E 1547.0−5408), & J1622−4950 (Camilo et al. 2006, 2007b; Levin et al. 2010)). On 25 April 2013, one day after the XRT aboard the *Swift* satellite detected flaring activity coincident with the Galactic Center (Degenaar et al. 2013), a short X-ray burst was observed by *Swift*/BAT showing characteristics similar to those usually observed from soft gamma-ray repeaters (Kennea et al. 2013c). Shortly thereafter, observations from the *NuSTAR* satellite identified the source as a magnetar with a $P_s = 3.76$ s spin period, and its radio pulsations were subsequently seen by the Effelsberg 100-m Telescope (Mori et al. 2013a,b; Eatough et al. 2013a). J1745−2900 was soon physically associated with the Galactic Center, located only ∼2.5" away from Sagittarius A* (Sgr A*) with a neutral hydrogen column density and dispersion measure (DM) consistent with being within ∼2 pc of the Milky Way's central black hole (Eatough et al. 2013b; Rea et al. 2013).

Early determinations of its spin-down $\dot{P}_s$ put J1745−2900 squarely within the

[2] `http://www.physics.mcgill.ca/~pulsar/magnetar/main.html`

magnetar population, having an inferred magnetic field strength at the equator $B_s \sim 3.2 \times 10^{19}$ G $\sqrt{P_s \dot{P}_s} \sim 1.6 \times 10^{14}$ G, a characteristic age $\tau_c \sim P_s/(2\dot{P}_s) \sim 9$ kyr, and a spin-down luminosity of $\dot{E} = L_{\rm rot} = 3.95 \times 10^{46}$ erg s$^{-1}(P_s^{-3}\dot{P}_s) \sim 4.9 \times 10^{33}$ erg s^{-1} (Rea et al. 2013). However, its estimated quiescent X-ray luminosity of $L_{\rm X,qui} < 10^{34}$ erg s^{-1} (Coti Zelati et al. 2015) may place J1745$-$2900 on the side of $L_{\rm X,qui}/L_{\rm rot} < 1$, opposite the "classic magnetars" but alongside the other three radio magnetars, high-B pulsars, and radio pulsars with X-ray emission (Rea et al. 2012).

Given the unique environment in which J1745$-$2900 resides, the detection of its radio pulses is somewhat surprising. Indeed, numerous surveys of the Galactic Center region covering $\sim$1–20 GHz have failed to find a pulsar within the central parsec (most recently, Johnston et al. 2006; Deneva et al. 2009; Macquart et al. 2010; Bates et al. 2011; Siemion et al. 2013). The discovery of this single magnetar has led to a windfall of implications for future discoveries (Chennamangalam & Lorimer 2014; Dexter & O'Leary 2014; Macquart & Kanekar 2015). Because of its proximity to the Galactic Center, J1745$-$2900 has the largest DM (1778 cm^{-3} pc) and rotation measure (-6.696×10^4 rad m^{-2}) of any known pulsar (Eatough et al. 2013b). The predicted value for the scattering timescale at 1 GHz, based on empirical relationships given its DM, is $\sim$ 1000 s (Krishnakumar et al. 2015; Lewandowski et al. 2015a), meaning that J1745$-$2900 would be undetectable at frequencies less than $\sim$5 GHz. The situation is exacerbated by the presence of an additional scattering screen in the Galactic Center (Cordes & Lazio 1997). Normally, the prospect of detecting distant radio pulsars above several GHz is bleak, since their average spectral index is $\sim$ -1.4 (Bates et al. 2013). However, because the other radio magnetars have flat/inverted spectra, one might expect to detect J1745$-$2900's unscattered pulse profile at high frequencies. In the analyses that follow, we will reiterate the finding that J1745$-$2900 has a significantly smaller scattering timescale than predicted (Spitler et al. 2014), and will show that J1745−2900 was much brighter at lower frequencies, having a very negative spectral index some 100 days after the onset of its outburst, even though more recent observations by Torne et al. (2015) showed the spectral index has since flattened.

In this paper, we analyze multi-frequency radio data over the first ∼100 days after J1745−2900's discovery, during which time there were two additional *Swift*/BAT-detected bursts on 7 June 2013 and 5 August 2013 (Kennea et al. 2013a,b). For two of our epochs, which bracket the third burst by ∼1 week on either side, we have simultaneous *Chandra* observations. These observations allow us to find the absolute alignment of the radio and X-ray profiles, and to look for correlated events. We comment on the spin evolution and timing, and examine the profile stability, the radio flux evolution, and the radio spectrum. Finally, we make global models of the profile evolution across the low frequency bands in order to examine the temporal and frequency dependencies of the scattering timescale and dispersion measure. We then discuss characteristics of this source in comparison with other radio-loud magnetars.

6.3 Observations

6.3.1 Radio

We made early detections of J1745−2900 during fourteen observing epochs with the 100-m Robert C. Byrd Green Bank Telescope (GBT) in three different frequency bands with various overlap: 1.1–1.9 GHz (5 epochs), 1.6–2.4 GHz (7 epochs), and 8.5–9.3 GHz (11 epochs) (PI: A. Possenti). Because each observation covers a large bandwidth, we refer to each set of data based on the IEEE radio band for which each of the receiver systems is named ("L-band", "S-band", or "X-band", respectively), instead of referring to specific (central) frequencies. Table 6.1 contains details of the observations. In all cases, we observed using the Green Bank Ultimate Pulsar Processing Instrument (GUPPI[3], DuPlain et al. 2008) in "incoherent search mode", recording dual-polarization time-series data in 2048 frequency channels with a temporal resolution of 0.65536 ms.

Each epoch's data were folded with the pulsar software library DSPSR[4] using a nominal ephemeris with a constant spin frequency (see §6.4.3) and the *Chandra*-determined position $\alpha_{\rm J2000.0} = 17^h45^m40\overset{s}{.}169$, $\delta_{\rm J2000.0} = 29°00'29\overset{''}{.}84$ (Rea et al. 2013). The data were initially folded into 1 min subintegrations, with 2048 profile phase bins across 128 frequency channels. We adopted the published dispersion measure value of 1778 cm^{-3} pc for averaging frequency channels together (Eatough et al. 2013b). Persistent, narrow-band radio frequency interference (RFI) was excised automatically; any remaining significantly corrupted channels or subintegrations were removed from the data by hand.

Calibration scans were taken for each observation using the local noise diode, pulsed at 25 Hz while on source. We recorded on- and off-source scans of a standard flux calibrator (QSO B1442+101) in each frequency band only during the final epoch (MJD 56516). We have used this one set of flux calibration scans to calibrate the whole data set. Standard programs from the PSRCHIVE[5] pulsar software library (Hotan et al. 2004; van Straten et al. 2012) were used to calibrate the absolute flux density scale of the noise diode, which is then used to determine the magnetar's flux density[6].

[3] `www.safe.nrao.edu/wiki/bin/view/CICADA/NGNPP`

[4] `http://dspsr.sourceforge.net/`

[5] `http://psrchive.sourceforge.net/`

[6] The PSRCHIVE calibration process produced unphysical results for the earliest S-band detection (MJD 56424); we have calibrated it by using an approximation based on the measured S-band

Table 6.1. J1745−2900: Summary of GBT Observations

UTC Epoch	MJD	Bands Observed	Approx. Length [min]
2013-05-04	56416	X	20
2013-05-12	56424	S*,X	122,200
2013-05-13	56425	X	60
2013-05-14	56426	X	49
2013-05-17	56429	S,X	70,53
2013-05-23	56435	X	50
2013-05-30	56442	X	58
2013-06-21	56464	X	54
2013-07-14	56487	X	71
2013-07-15	56488	L,S	120,132
2013-07-27	**56500**	L,S,X	186,108,68
2013-07-28	56501	L,S*	133,117
2013-08-03	56507	L,S	112,75
2013-08-12	**56516**	L,S,X	120,60,56

Note. — The listed dates and MJDs for the epochs are representative of the majority of the epoch, not the start time; observations on the same day were taken in tandem. The two boldfaced epochs are those for which we have simultaneous observations with *Chandra*. The lower half (400 MHz) of the two S-band observations with an asterisk were corrupted and unusable. The horizontal lines separate the epochs during which the three observed types of X-band profile are seen (see §6.4.2 and Figure 6.2).

The combination of the large amount of observed scattering (§6.4.5), the pulsar's spectrum (§6.4.4, Figure 6.7), receiver roll-off, and the presence of gain variations (see below) rendered significant portions of the ends of L-band useless. Namely, there was no pulsed signal in the lower 300 MHz portion of L-band, which we masked from further analysis, along with the top 50 MHz (which is part of the overlap with S-band). In combination with the narrow-band RFI, this left less than ∼400 MHz of clean, usable bandwidth. Similarly, at S-band we had to remove the lower ∼100 MHz and the upper ∼25 MHz, and in total ∼625 MHz of usable band remained[7]. Only 3% of the data was clipped from either end of X-band, with a total of 10% removed. We took these seemingly draconian measures to offset the original data quality and to ensure that the time- and frequency-averaged profiles were of reasonably high quality (e.g., see Figure 6.1). This was enabled by the source's relatively large flux density.

The data quality situation at X-band was still more complicated. As also noted by Lynch et al. (2014) in their investigation of this magnetar, large gain variations on timescales from a fraction of a pulse period to several seconds (visible in the time-series data) are prevalent in X-band at the GBT, when pointed at the Galactic Center. The variations did not (necessarily) integrate away over hour-long observations and are representative of a stochastic red-noise process. We attribute these variations to changes in atmospheric opacity (Lynch et al. 2014) and/or small pointing errors, noting a strong resonance in the GBT X-band pointing very near 0.3 Hz[8]. The gain variations would be manifested by the relatively small beam of X-band (∼1.4′,

---

system equivalent flux density and the radiometer equation (Lorimer & Kramer 2005, cf. §7.3.2, Equation 7.12). The result is reasonable, given that the next S-band observation five days later has a comparable flux density (see Figure 6.6).

[7] In two epochs, however, instrument problems left only half of S-band viable. See Table 6.1.

[8] Even though the average pointing errors at X-band are only on the order of several arcseconds at mid-elevations and mild wind conditions, the power spectrum in elevation offset shows resonances overlapping with the magnetar's spin frequency (0.27... Hz). See `http://www.gb.nrao.edu/~rmaddale/GBT/Commissioning/Pointing_Gregorian_HighFreq/PntStabilityXBand.pdf` for details.

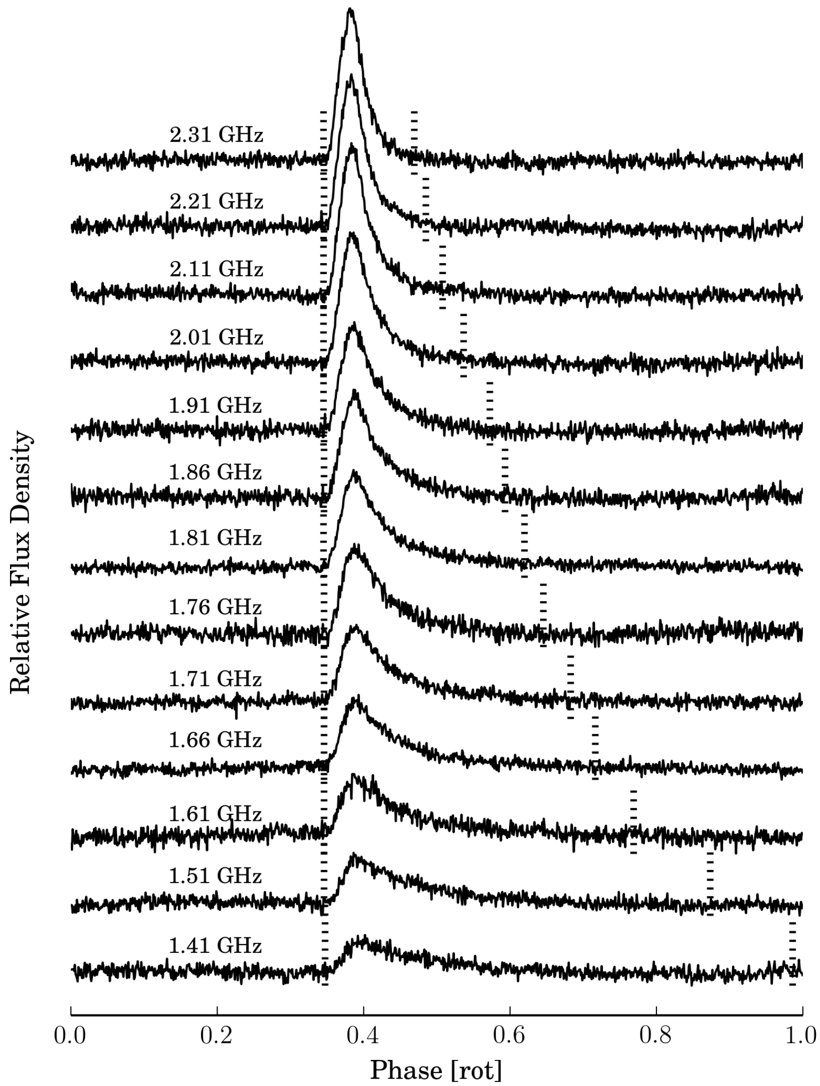

Fig. 6.1 – Examples of L- and S- band profiles averaged over all epochs. The profiles are shown with 1024 phase bins for clarity. These data are aligned via a wideband portrait model, as described in §6.4.5. In general, the un-averaged profiles were also of good quality, with only minor systematics in the baseline. The total bandwidth covered across these two bands is about 1 GHz, from ∼1.4 to 2.4 GHz; 25 MHz of data were averaged for each of these profiles, with their center frequencies shown. The profiles were very well described by a single scattered Gaussian component, and so we do not over-plot the wideband model. The vertical dotted lines show examples of on-pulse regions used for the flux density measurements. See §6.4.4 for details.

compared to $\sim 6'$ and $\sim 9'$ for S- and L-band) oscillating over the crowded, bright Galactic Center (the central parsec extends $\sim 0.4'$, and the separation of J1745−2900 from Sgr A* is only $\sim 0.04'$ (Rea et al. 2013)). Additionally, it is likely that the baseline variations are much less prominent at low frequencies because they act as "zero-DM" signals that get smeared out when the pulsar's signal is dedispersed. Lynch et al. (2014) also state that the effect may be a function of elevation, which fits with our pointing-resonance hypothesis, since the influence of variable elements like the wind will be a function of elevation. The persistence and variability of these variations can be seen in Figure 6.2.

The analyses that follow utilized these folded profiles in a variety of reduced forms. Unless otherwise noted, the reduced radio data have 2048 profile bins ($\sim$7.2 ms per bin), 32 frequency channels (25 MHz per channel), and 5 min subintegrations; in this work, we only consider the total intensity profiles.

6.3.2 X-ray

During two of our radio epochs, MJD 56500 and MJD 56516, we obtained simultaneous observations of J1745−2900 with the *Chandra* X-ray Observatory (Obs. IDs 15041 & 15042; PI: D. Haggard). Table 6.2 contains details of the X-ray observations (for further details see Coti Zelati et al. 2015). The field of the first observation is shown in Figure 6.3; the second observation was essentially the same. In each observation, J1745−2900 was positioned on the back-illuminated chip S3 of the Advanced CCD Imaging Spectrometer (ACIS, Garmire et al. 2003) instrument. The data were reprocessed with the *Chandra* Interactive Analysis of Observations software package (CIAO, version 4.6, Fruscione et al. 2006) and the calibration files in the CALDB release 4.5.9.

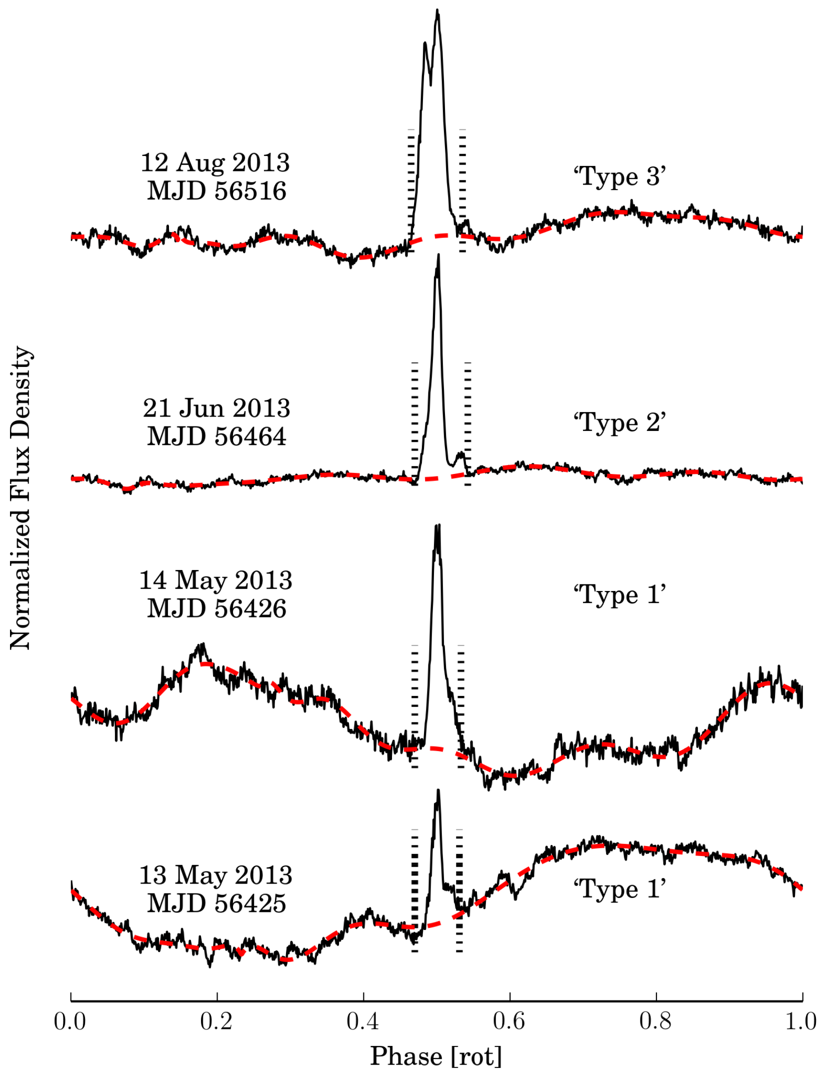

Fig. 6.2 – Examples of time- and frequency-averaged X-band profiles. The profiles are shown with 1024 phase bins for clarity. The baseline variations were removed on a profile-to-profile basis by fitting a high degree polynomial (red dashed lines) to the off-pulse region (outside the dotted lines) in order to make measurements of the flux density (see §6.4.4). The on-pulse phase window varied in size between about 6 and 8%. The profile evolved monotonically from one "type" to the next (see §6.4.2 and Table 6.1).

Table 6.2. J1745−2900: Summary of Simultaneous *Chandra* Observations

Obs. ID	Radio epoch [MJD]	Exposure time [ks]	Net source counts [10^3]	RMS pulsed fraction [%]
15041	56500	45.4	15.7	28.8 ± 1.5
15042	56516	45.7	14.4	28.9 ± 1.8

Note. — The 1σ uncertainties for the RMS pulsed fractions were determined from Monte Carlo simulations (cf. Gotthelf et al. (1999)). By another measure, the pulsed fractions — defined as the difference between the profile maximum and minimum divided by their sum — are ∼48%. The folded profiles are shown in Figure 6.5.

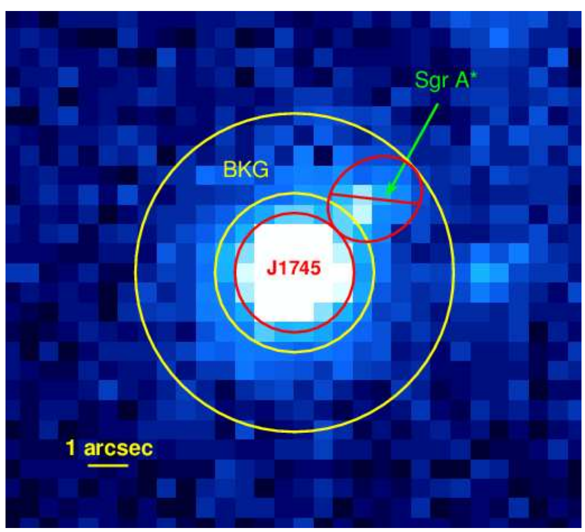

Fig. 6.3 – *Chandra* field of J1745−2900 for observation 15041. 1 ACIS pixel = 0.492″. The source counts were taken from the central-most encircled region (red circle). Background counts were extracted from the annulus between the outer two (yellow) circles, excluding the area marked as “Sgr A*”. We account for pile-up as described in §6.3.2.

In both observations, J1745−2900 was bright enough to cause pile-up in the ACIS detector. A "pile-up map" created with the CIAO tool `pileup_map` confirmed that mild pile-up was present. Exclusion of data near the center of the point-spread function (PSF) from the analysis would have resulted in the loss of too many photons (63% of the source counts were in the two central pixels). Moreover, the external part of the PSF contained a substantial number of counts from Sgr A*. We thus decided to proceed as follows.

We extracted the source counts from a circular region centred on J1745−2900 with a $1.5''$ radius (see Figure 6.3); this region includes the piled-up events. This area covers ∼85% of the *Chandra* PSF (encircled energy fraction) at 4.5 keV. A larger radius of 2–$2.5''$ would let in more counts from Sgr A* and would only marginally increase the encircled energy fraction. Because of the complex environment, the background spectrum needed to be extracted close to the source. We used a thin annulus (with radii of $2''$ and $4''$), excluding a bright area associated with Sgr A*. The spectra, the ancillary response files and the spectral redistribution matrices were created using `specextract`. Following Rea et al. (2013), we adopt a pure blackbody for the spectral shape. We corrected the spectra using the pile-up model by Davis (2001), as implemented in the modeling and fitting package SHERPA (Freeman et al. 2001). The pile-up fraction, estimated by fitting the `jdpileup` model, is 3.7% for the first observation, and 4.1% for the second. We did not attempt any correction of the light curves; the pile-up fraction is modest and, in general, pile-up affects spectra more than it does light curves and pulse profiles[9]. The spectral model fits were acceptable only when the pile-up model component was included. A summary of the spectral fits is given in Table 6.3.

[9] This is true unless the pulse profiles are strongly dependent on energy, which is not the case for J1745−2900, though we refer the reader to Coti Zelati et al. (2015) for further details.

Table 6.3. J1745−2900: *Chandra* Spectral Results

Obs. ID	μ[a]	f[a]	$N_{\rm H}$ [10^{23} cm^{-2}]	kT[b] [keV]	R[b] [km]	Observed flux[c] [10^{-12} erg cm^{-2} s^{-1}]	Luminosity[c] [10^{35} erg s^{-1}]	$\chi^2_{\rm red}$ (dof)
15041	$0.50^{+0.31}_{-0.05}$	99.8%	1.26 ± 0.03	0.82 ± 0.01	$2.34^{+0.13}_{-0.17}$	8.9 ± 1.2	$3.2^{+0.4}_{-0.3}$	1.00 (288)
15042	$0.48^{+0.25}_{-0.08}$	97.1%	1.23 ± 0.03	0.83 ± 0.01	$2.16^{+0.13}_{-0.16}$	$8.1^{+1.4}_{-1.1}$	2.8 ± 0.4	1.00 (287)

[a]Parameters of the `jdpileup` SHERPA pile-up model; μ is the grade-migration parameter and f is the fraction of the PSF treated for pile-up, required to be in the range 85–100%. For details, see Davis (2001) and "The *Chandra* ABC Guide to Pileup".

[b]The blackbody temperature and radius are calculated at infinity and assuming $D = 8.3$ kpc (Genzel et al. 2010), which is assumed throughout this work.

[c]In the 0.3–8 keV energy range; for the luminosity we again assumed $D = 8.3$ kpc.

Note. — The abundances used in the absorbed blackbody model are those of Anders & Grevesse (1989) and photoelectric absorption cross-sections are from Balucinska-Church & McCammon (1992). See Coti Zelati et al. (2015) for a complete treatment of these observations in the context of a long-term X-ray monitoring campaign. Parameter uncertainties in the table are 1σ.

6.4 Results

6.4.1 Transient Events

J1745−2900 is known to show narrow individual pulses (Spitler et al. 2014; Bower et al. 2014; Lynch et al. 2014), similar to the radio magnetar J1622−4950 (Levin et al. 2012). We performed a cursory analysis of J1745−2900's individual pulses in our X-band data, seeking only to find anomalous burst-like events in the radio data that might be coincident or correlated with X-ray features or flares. For this, we took two approaches. In the first case, we folded the raw data into single-rotation integrations, approximately maintaining the original temporal resolution, averaging over frequency, and summing the polarizations. These data were inspected visually. In the second case, we analyzed the raw data with the Presto[10] pulsar software package. Here, we applied an RFI mask to the raw data with `rfifind`. We then made a dedispersed[11], frequency-averaged time-series with `prepdata` for each X-band epoch, and searched for single pulses with the boxcar-convolution algorithm implemented in `single_pulse_search.py`. We repeated this process on the unmasked raw data.

Single pulses were detected; indeed, one to several pulses are visible by eye during almost every rotation. However, we saw no anomalously large single pulses or other bursts. The distributions of estimated single pulse energies all peak at $\lesssim$1 times the average profile energy and were inconsistent with power-law distributions. The phases of the single pulse arrival times were consistent with occuring within the on-pulse window, and the distributions of the resolved single pulse widths peaked near 3-4 samples $\approx$ 2 ms, in agreement with the X-band scattering timescales found by

[10] `http://www.cv.nrao.edu/~sransom/presto/`

[11] We used the Eatough et al. (2013b) DM of 1778 cm^{-3} pc, and compared the results to those from times-series dedispersed at 0 cm^{-3} pc and twice the nominal DM in order to discriminate between transient RFI and candidate pulses.

Bower et al. (2014).

Similarly, the (unfolded) X-ray light curves during the two simultaneous observations, binned from 0.5 to 5000 s, were featureless and constant. The χ^2 probability of constancy was high for both observations, regardless of the choice of binning (>30% and frequently approaching 100%). Due to the uniformly poor quality of the X-band data as previously described, we refrain from further analysis or discussion of this aspect of J1745−2900 and direct the reader to the observations of its X-band single pulses as observed with the Very Large Array and the GBT in Bower et al. (2014) and Lynch et al. (2014), respectively.

6.4.2 Profile Variability

Figure 6.2 shows examples of the three general types of observed X-band profiles, as well as corresponding examples of our baseline removal and on-pulse determination. The transition between "Type 1", with a single main component having a trailing-side shoulder and a more quickly rising leading edge, and "Type 2", with the main component having a leading-side shoulder and a nub feature on the trailing side, happens more than three weeks after the X-ray burst on MJD 56407 and more than two weeks before the burst on MJD 56450. The "Type 1" shape was seen as early as a week after the discovery (Eatough et al. 2013a) and published in Eatough et al. (2013b). Similarly, the transition between "Type 2" and "Type 3", which has a larger two-peaked component, happened more than two weeks after the burst on MJD 56450 and more than three weeks before the burst on MJD 56509. For these reasons, we do not associate the profile types (which are most likely not absolutely discretized) with the observed X-ray bursts. Within a single observation, the profile shape did not change between 5 min subintegrations.

Lynch et al. (2014) also documented the time-variability of J1745−2900's X-band profile as seen with the GBT. As their first observation is coincident with our last observation, they have also seen the "Type 3" shape, which persists and evolves during most of what they have labeled a "stable state". This "stable state" is characterized by relatively smooth profile transitions, a gradual flux evolution, and a phase-connected timing solution — all in contrast to what they call an "erratic state", which is onset sometime after MJD 56682. Later in their observations, during epochs with MJDs 56794 and 56865 (both in the "erratic-state"), they see a profile resembling what we have labeled "Type 1". We note that we did not witness any of the very sporadic profile variability seen in Lynch et al. (2014) associated with the "erratic state" (e.g., the drastic profile changes seen in their last two observations, separated by only eleven days), but rather we observed each of these three types only for a single interval of time.

6.4.3 Timing

Between having bursts, glitches, unstable profiles, and timing noise, magnetars are notoriously some of the hardest pulsars to time (cf. the original radio magnetar J1809−1943 (Camilo et al. 2006), or see a recent review of magnetars in Mereghetti (2013)). As is evident from the X-ray and radio timing in Coti Zelati et al. (2015), Lynch et al. (2014), and Kaspi et al. (2014), obtaining a single phase-connected timing solution for J1745−2900 is difficult, due to a significant level of timing noise. Here, we measure an overall average spin-down for the purpose of summing the data in each epoch.

Pulse times-of-arrival (TOAs) were measured by cross-correlating the time- and frequency-averaged data profiles with smoothed, "noise-free" template profiles using standard PSRCHIVE routines. The templates are generated by arbitrarily aligning and averaging all of the data for which the template is used. Single templates were used for the L- and S-band data, but three separate templates were used for X-band, depending on the profile observed, as discussed in §6.4.2. Arbitrary phase offsets were fit between TOAs measured from all of the different templates as part of the timing models. These phase offsets serve to align the template profiles, but do so indiscriminately with respect to pulse broadening from interstellar scattering; this has the effect of biasing DM estimates if one tries also to measure the dispersive delay between TOAs of different frequencies. See §6.4.5 for our DM measurements based on wideband modeling of the L- and S-band data.

Figure 6.4 shows the measured values of the spin frequency f as a function of time. The average measured spin-down of $\dot{f}_{avg} = -8.3(2) \times 10^{-13}$ Hz s^{-1} was sufficient to average the data in each epoch with negligible smearing for the flux measurements (§6.4.4), and is a reasonable approximation for the overall trend in the spin evolution[12]. This average value also lies between the two $\dot{f}$ values presented in Table 2 of Kaspi et al. (2014) for the same range of dates.

Although we are not interested in a full timing solution for these data in this work, we found corroborative results when following the suggestion in Kaspi et al. (2014) that there is an abrupt change in $\dot{f}$ around the time of the *Swift*/BAT-observed X-ray burst on MJD 56450. Namely, while a single, predictive timing solution was not found, our pre- and post-burst TOAs are described by two simple phase-coherent solutions with parameters $\dot{f}_{\rm pre} = -5.005(1) \times 10^{-13}$ Hz s^{-1}, $\dot{f}_{\rm post} = -9.4799(5) \times 10^{-13}$ Hz s^{-1}, and $\ddot{f}_{\rm post} = -2.696(6) \times 10^{-20}$ Hz s^{-2}. These values are in good agreement with those in Coti Zelati et al. (2015), Kaspi et al. (2014), and Rea et al. (2013), although

[12]Quantities in parentheses represent the 1σ uncertainty on the last digit in the respective measurement throughout the paper.

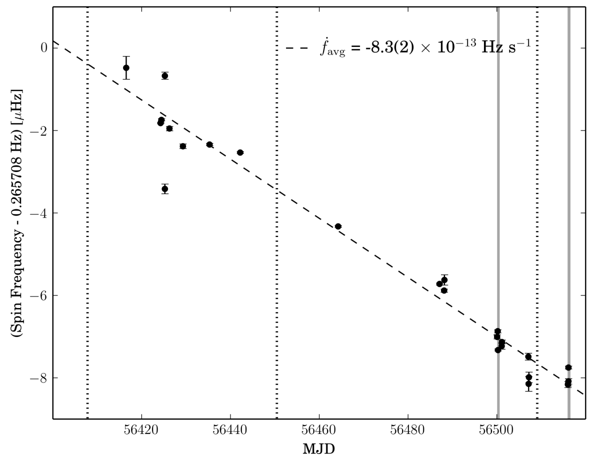

Fig. 6.4 – Average spin evolution of J1745−2900. The three vertical dotted lines correspond to the three X-ray bursts detected by *Swift*/BAT. The two vertical grey bars cover our *Chandra* observations. Measurements from the two early S-band observations are not included, nor from the X-band epoch on MJD 56425, as they were very significant outliers. The quoted uncertainty does not include the residual scatter.

we were not sensitive to $\ddot{f}_{\mathrm{pre}}$. We could only obtain a single phase-connected timing solution for all of the TOAs by using five spin frequency derivatives, which is not a predictive ephemeris.

Profile Alignment

In Figure 6.5, we present the absolute alignment between the *Chandra* 0.3–8 keV X-ray profiles and the GBT radio profiles in L-, S-, and X-bands. We determined an independent ephemeris for each of the two epochs from the radio data by fitting TOAs from each day for only the spin frequency, fixing the spin-down parameter at the average value reported above. These TOAs were measured from the frequency-averaged data with 5 min subintegration resolution. The phase-zero time was referenced to the arrival of infinite-frequency radiation at the Solar System barycenter, which assumes a constant dispersion measure of 1778 cm^{-3} pc between the two observations. The X-ray photon arrival times were barycentered also using the sky position given in §6.3.1 and the JPL Planetary Ephemeris DE-405. These events were folded into pulse profiles with 64 phase bins using the corresponding epoch-specific ephemeris by the `prepfold` program of PRESTO. The alignment based on folding using a single ephemeris for both epochs — either the post-burst or the multiple frequency derivative ephemeris — yielded indistinguishable results. This is reasonable, since the RMS timing residual from either of those ephemerides is on the level of individual bins. On the other hand, it may be surprising that there seemed to be no interruption in the "post-burst" ephemeris; the third detected X-ray burst occurred at the midpoint between the two simultaneous radio/X-ray epochs.

We modeled each of the two X-ray profiles with four Gaussian components to measure the relative offsets with respect to the radio profiles. The offsets and their uncertainties were determined from Monte Carlo trials, where "offset" here refers to the phase that maximizes a cross-correlation such as the one prescribed in Taylor (1992). There was a small offset between the X-ray models, $\lesssim$0.02 rot. A difference in DM would shift the relative phase between the X-ray profile and the S-band profile (our fiducial profile) only by $\sim 3\times10^{-4}$ rot per unit DM [$\mathrm{cm^{-3}}$ pc]. Even for the DM difference of $\sim$17 $\mathrm{cm^{-3}}$ pc measured between these epochs (see §6.4.5 and Figure 6.8), the phase difference is $\sim$0.005 rot. The remaining offset can be explained by a combination of the variability of the X-ray profile and timing noise, with the former being dominant. After removing this difference, the offsets with respect to the radio profiles do not change between the two days within the variance of the measurements. The phase offset relative to the S-band profile is approximately 0.15(1) rot. The radio magnetars J1809−1943 and J1550−5418 both also show rough alignment of pulsed radio emission with their X-ray profiles (Camilo et al. 2007a; Halpern et al. 2008), whereas no pulsed X-ray emission has been detected from J1622−4950 (Anderson

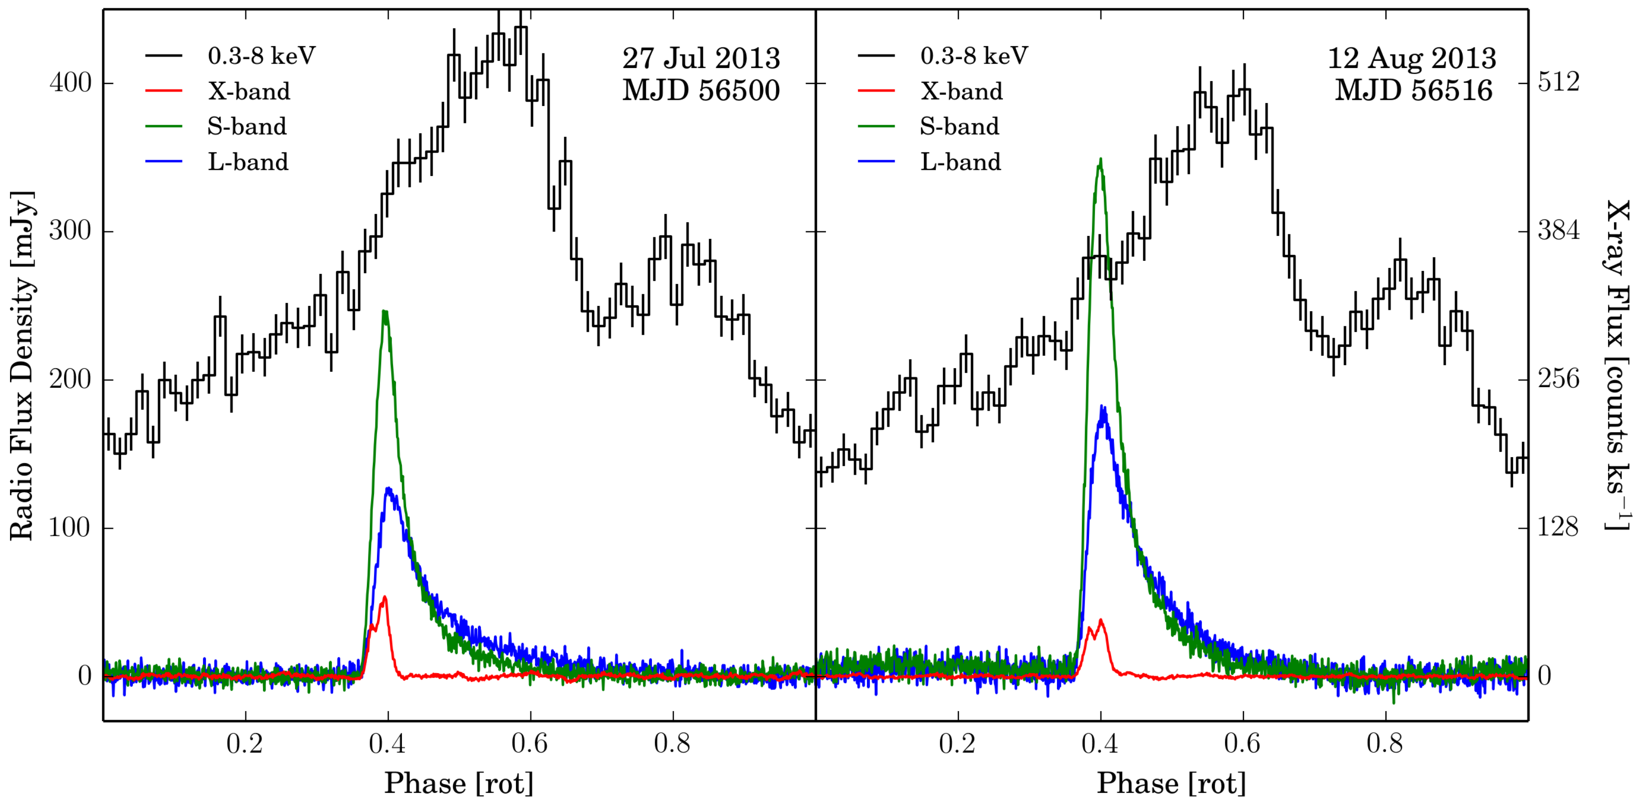

Fig. 6.5 – Absolute phase alignment of J1745−2900's radio and X-ray profiles determined separately on two days. Note that the brightest radio profile is seen in S-band (see Figure 6.7). The profiles have 1024 and 64 phase bins, respectively, and are shown as they would be observed at the Solar System barycenter for phase-zero MJDs 56499.98000761 and 56515.96999979, referenced to infinite frequency. The assumed dispersion measure is 1778 cm^{-3} pc. During two later *XMM-Newton* observations (presented in Coti Zelati et al. (2015)), there is a peculiar, narrow feature seen in the otherwise broad X-ray profile near the phase of radio emission as shown here (see text).

et al. 2012).

The two double-peaked X-ray profiles appear essentially featureless. The RMS pulsed fractions are given in Table 6.2. There are not sufficient data to decompose the profiles into energy bands to look for meaningful spectral dependencies, although we wish to point out a possible transient feature that appears in the *XMM-Newton* data recently published by Coti Zelati et al. (2015). In the energy-dependent *XMM-Newton* profiles of Figure 4 from Coti Zelati et al. (2015), there is a conspicuous narrow feature on the leading edge of the double-humped X-ray profile that is close to the phase of radio emission (within ∼0.05 rot). It appears most prominently around phase 0.55 in the 0.3–3.5 keV profile of the third *XMM-Newton* observation (with Obs. ID 0724210501). It is also seen in two of the other three energy-dependent profiles (except for the highest energy 6.5–10.0 keV profile), contributing to the integrated flux in the energy-averaged profile. A similar feature is seen at the same phase in the first *XMM-Newton* observation (with Obs. ID 0724210201) to a lesser extent. According to their table, these observations were separated by 23 days, with the first occurring 19 days after the *Chandra* observations presented here (which are also included in Coti Zelati et al. (2015)). The three *Chandra* observations and the one *XMM-Newton* observation taken during these 23 days show no obvious feature, despite covering the same range of energies, although *Chandra* recorded only between 10 and 50% of the counts as did by *XMM-Newton*. Therefore, without additional observations, it remains only a peculiarity.

6.4.4 Radio Flux Density

From the radio data, we made measurements of J1745−2900's flux density as a function of time and frequency. We measured the mean flux densities in 50 MHz wide channels and used a weighted average of these measurements to obtain representative flux densities for each band, per epoch.

For all of the L- and S-band profiles, we defined "on-pulse" regions as follows. A model pulse profile for each frequency was determined from the wideband modeling described in §6.4.5. We then found the smallest range of pulse phases that contained 99% of the integrated flux density of the model profile. Examples of the on-pulse windows for the scattered L- and S-band profiles can be seen in Figure 6.1. The mean flux density was calculated by averaging the observed flux density in the window and scaling it by the duty cycle. The uncertainties were estimated by measuring the mean noise level in the last quarter of each profile's power spectrum[13]. We accounted for systematics in the residual profile by adding the scaled, residual mean flux density to the uncertainty in quadrature. These corrections were small, as the reduced χ^2 values of the residuals were usually <1.5 and always <2.0.

The measurement of the X-band flux densities was complicated by the dynamic baseline variations mentioned in §6.3.1, as well as the intrinsic variability of the profile shape. We used polynomial functions to remove the baseline variations on a profile-to-profile basis (e.g., see Figure 6.2). For these profiles, we first centered each profile to be near phase 0.5 to avoid edge-effects of the polynomial fit from affecting the on-pulse region. A high degree polynomial function was fit to the baseline of each profile, where in the first iteration an on-pulse window with a duty cycle of 6% was blanked out from the fit to avoid initially over-estimating the noise[14]. The level of the residual off-pulse noise was calculated, and then the on-pulse window was

[13]This is a robust method to estimate the off-pulse variance, assuming the profile is resolved (e.g., see Demorest 2007).

[14]None of the profiles had a smaller duty cycle than 6% and a polynomial of degree 15 was used; this was the smallest degree polynomial that reasonably and automatically removed systematic baseline trends from all of the profiles without having to also vary the degree of the polynomial on a profile-to-profile basis.

widened until the flux density at the edges of the on-pulse region dropped below the noise level. The baseline polynomial was then refit to the original profile, but with the new on-pulse window blanked out. The mean flux density and its uncertainty were calculated in these baseline-removed, on-pulse windows as described for the lower frequency data above, but a systematic error was added in quadrature to the uncertainty that represented the mean flux density across the on-pulse phase window removed by the polynomial fit. This tested method gives dependable, conservatively estimated X-band flux densities.

Flux Evolution

The radio flux evolution of J1745$-$2900 is shown in Figure 6.6. The mean X-band flux density increases rapidly in the first half of our observations, increasing by at least a factor of $\sim$6 over fifty days, and then tapers off at the 1 mJy level. The earliest reported measurement of J1745$-$2900's X-band flux density was $\sim$0.2 mJy, taken with the Effelsberg 100-m Radio Telescope, consistent with our GBT measurement two days later (Eatough et al. 2013a). Our data show a similar increase in the low frequency flux densities. The S-band flux increases by about an order of magnitude over ninety days, and in our last five observations covering about thirty days, the average L- and S-band fluxes increase by a factor of two. Given the measured scattering timescales for J1745$-$2900 (see §6.4.5) and the recently measured proper motion of the pulsar, the timescale for refractive scintillation to be important is much larger than the span of our observations (see Bower et al. (2015) for further discussion).

Having picked up where we left off, Lynch et al. (2014) increased the cadence of GBT X-band observations after MJD 56516 and found a similar, slow increase

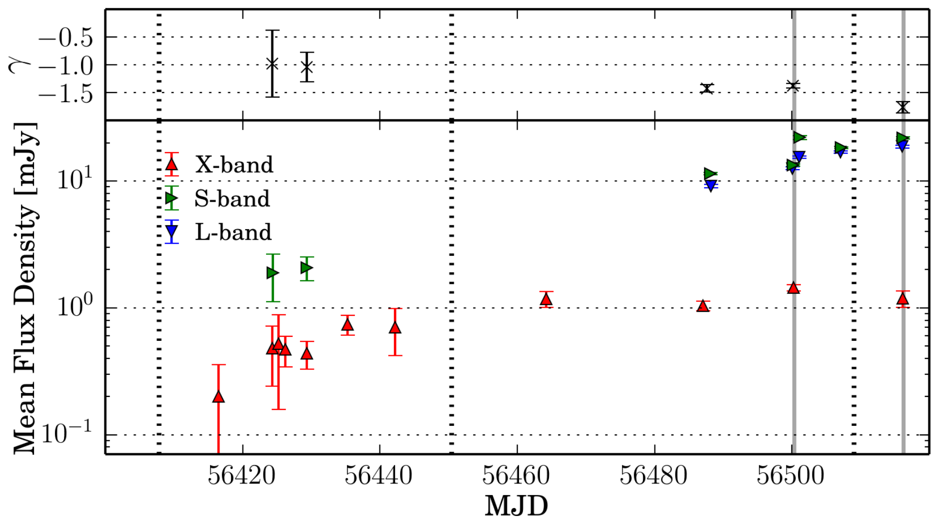

Fig. 6.6 – The early radio flux (bottom panel) and spectral (top panel) evolution of J1745−2900 over 100 days from the observations in Table 6.1. The vertical demarcations are the same as in Figure 6.4. Lynch et al. (2014) find a continuation of the slow, steady increase in X-band flux for another six months, which is followed by what they call an "erratic state". The apparent excess average S-band flux density during MJD 56501 is explained by the fact that the lower half of the band was corrupted (see Table 6.1), and the pulsar's flux density apparently increases with frequency in this range (see Figure 6.7). The average value of the spectral index γ is about −1.4; see text for details.

of the flux, up to ∼3 mJy, over the next 170 days. As already mentioned, after this "stable state" of slow, steady flux increase, the authors found that J1745−2900 entered an "erratic state", characterized in part by a larger and highly variable X-band flux, similar to what was seen in two other radio magnetars (Camilo et al. 2007a; Levin et al. 2012). Superimposed on top of this radio flux evolution is a relatively slow decay of the X-ray flux, compared to other magnetars (Rea et al. 2013; Kaspi et al. 2014; Lynch et al. 2014; Coti Zelati et al. 2015). Between our two simultaneous GBT/*Chandra* observations separated by ∼15 days, the radio flux increased by ∼60% while the X-ray flux decreased by ∼10%. This trend (seen here and in Lynch et al. (2014)) is opposite to those of the other radio magnetars, which show decreasing radio and X-ray flux with time over the course of an outburst (Rea et al. 2012).

Radio Spectral Index

Because we have essentially simultaneous observations[15] of J1745−2900 in frequency bands spaced by two octaves, we can measure the spectral index γ, where $S_\nu \propto \nu^\gamma$ for flux density S_ν at frequency ν. The upper panel of Figure 6.6 shows γ as measured between the average X-band flux density and the combined average flux densities of the lower frequency band(s). The error bars were approximated by varying the average fluxes within their measurement uncertainties. The decorrelation bandwidth for diffractive scintillation is much smaller than even our native frequency resolution and will not be a source of variability here.

There is no large, obvious stochasticity, as opposed to, for example, J1809−1943 (Lazaridis et al. 2008), but there may be a trend. Shannon & Johnston (2013) report two early measurements of γ across the bands spanning 4.5–8.5 GHz and 16–20 GHz. The first measurement on MJD 56413 is close to -1.0 in the high frequency band, though it is closer to 0.0 in the lower frequencies, and the second on MJD 56443 is ~ -1.0 across both bands, consistent with our measurements more than two weeks prior. Our three later measurements indicate a significantly steeper spectrum. The average value for γ of -1.4 is tantamount to the average spectral index for normal pulsars across gigahertz frequencies as reported in Bates et al. (2013). Camilo et al. (2007c) and Anderson et al. (2012) both make mention of a general steepening of the spectral indices of J1809−1943 and J1622−4950, respectively, despite remaining much flatter than what is seen in J1745−2900. However, (Lazaridis et al. 2008) finds the opposite for J1809−1943 in later observations.

This finding apparently breaks the mold set by the other three radio magnetars, which have essentially flat (or inverted) spectra (Camilo et al. 2006, 2008; Levin et al.

[15] In one case, the X-band observation was taken a day earlier; see Table 6.1.

2010; Keith et al. 2011). However, no firm conclusions can be drawn from this handful of measurements from early times in J1745−2900's outburst, especially knowing that the other radio magnetars also show a variable radio spectrum (Camilo et al. 2007c; Lazaridis et al. 2008; Anderson et al. 2012). In fact, at the time of writing, the findings of Torne et al. (2015) suggest that at much later times (a year after the present observations), the radio spectrum of J1745−2900 between 2 and 200 GHz was much flatter, with $\gamma = -0.4(1)$.

Spectral Shape

One example of J1745−2900's radio spectrum is shown in Figure 6.7; the spectra from the other days are qualitatively similar. The spectrum shows a non-power-law increase in flux between 1.4 and 2.4 GHz, with a possible peak near 2 GHz. The inverted log-parabolic shape is reminiscent of what have been called "gigahertz-peaked spectra" (GPS) pulsars (Kijak et al. 2011, 2013; Dembska et al. 2014, 2015), although the GPS pulsars supposedly have a much broader spectral shape, over a dex in frequency. For reference, we fit a log-parabola to the low frequency points, the parameters of which are given in the figure.

It is difficult to explain the spectral shape we see in the lower frequencies. It is conceivable that the dense, unique environment near J1745−2900 in the Galactic Center significantly alters the spectral shape of radio emission between 1 and 10 GHz (e.g., via free-free absorption, although the detection of Sgr A* at 330 MHz implies a low free-free optical depth of $\lesssim$1 (Nord et al. 2004)), but it is difficult to draw any conclusions without a dedicated set of observations.

Another possibility is that we have systematically under-estimated the flux: one well known source of bias comes from under-estimating the flux at low frequencies due

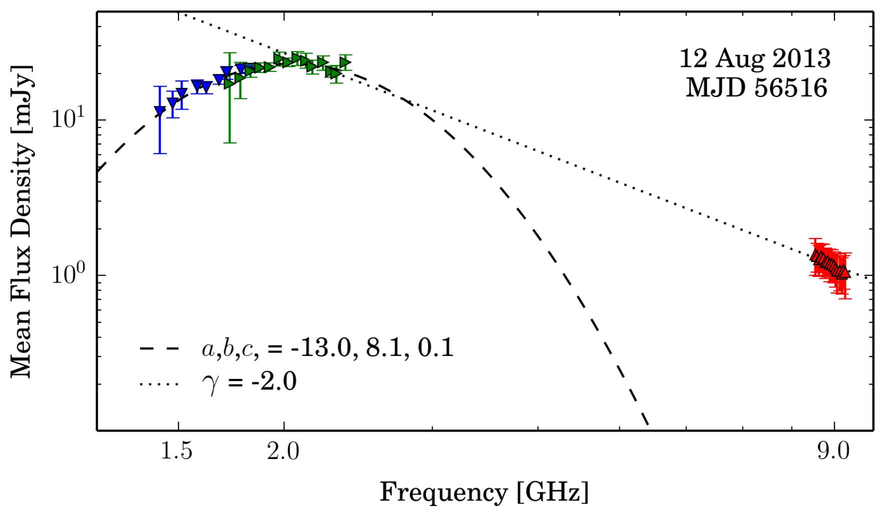

Fig. 6.7 – An example of J1745−2900's radio spectrum from the brightest observed epoch, MJD 56516. The markers are as in Figure 6.6; note that the flux densities agree in the ∼100 MHz overlap between L- and S- band. A similar inverted parabolic shape over log-frequency is seen during the other sets of (nearly) simultaneous observations, which is reminiscent of the so-called GPS pulsars. The coefficients a, b, and c of the fitted dashed parabola ($\log_{10}(\mathrm{S}_\nu) = ax^2 + bx + c$, for x = $\log_{10}(\nu)$) are given in the plot, along with the spectral index γ, which for this plot was fitted between the *peak* of the parabola and the X-band data.

to significant area in the scattering tails being lost in the calculation of the baseline flux. However, even at 1.4 GHz the scattering timescale is ∼500 ms ≈ 0.13 rot (see §6.4.5). In the worst case of a Kolmogorov scattering index (−4.4), the scattering timescale at our lowest frequency is no more than ∼20% of a rotation. As mentioned in Kijak et al. (2011) and treated graphically in Macquart et al. (2010), the pulsed fraction drops by only ∼10% when the scattering timescale is *half* the pulse period. Therefore, we can suggest that at worst we are underestimating the L-band flux densities at the ∼10% level, but this still would imply a positive or approximately flat spectral index between 1.4 and 2.4 GHz; the observed flux density drops precipitously somewhere thereafter.

A more promising, albeit provisional possibility has been offered up by recent modeling of the Shannon & Johnston (2013) observations. Lewandowski et al. (2015b) make a case study of J1745−2900 to demonstrate the possibility of thermal free-free absorption as the explanation for the GPS. For J1745−2900, the authors suggest a combination of an expanding ejecta and/or an external absorber to explain the changing spectrum seen early after the initial outburst in Shannon & Johnston (2013). The free-free absorbed model spectra offer a reasonable explanation for the lack of low-frequency detections of J1745−2900 immediately after the initial outburst and detections above 4 GHz; our two early S-band observations may support this idea. Our spectra from three months later may also inform the story of an evolving or endemic free-free absorbing medium in the environment of J1745−2900.

6.4.5 Wideband Portrait Model

As is evident from Figure 6.1, J1745−2900 has a highly scattered, simple profile across a gigahertz bandwidth, from 1.4 to 2.4 GHz. For a nominal DM value of 1778 cm^{-3} pc, there is a delay of $\sim$0.66 rotations across this band, which is easily measurable. All of the average L- and S-band profiles showed prominent scattering tails from multipath propagation through the interstellar medium (ISM). The quality of the data permitted us to make "wideband" measurements of both the DM and the scattering timescale τ, as well as its power-law index α, on an epoch-to-epoch basis[16].

For this, we used the methods and augmented software described in Pennucci et al. (2014) to make a wideband "portrait"[17] model for each of the five epochs where we have both L- and S-band observations. For each of these epochs we combined the data from the two low frequency bands in a fit for a global portrait model that included a single scattered Gaussian component with profile evolution parameters, a

[16]The two earliest S-band observations were exceptions; corrupted data, low signal-to-noise ratios, and the lack of L-band data resulted in uninformative measurements of the DM, τ, and α. This is also the reason these observations were excluded from the average $\dot{f}$ measurement in §6.4.3.

[17]We use the word "portrait" to mean the total intensity profile as a function of frequency.

constant baseline term, a phase offset between the bands, DMs for each band, and the scattering index α. The scattering timescale is defined in the usual way by assuming a one-sided exponential pulse broadening function for the ISM, so that an observed profile $p(\varphi)$ is the convolution given by

$$p(\varphi) = g(\varphi) * e^{-\frac{\varphi P_s}{\tau}} \mathrm{H}(\varphi), \tag{6.1}$$

where φ is the rotational phase, P_s is the spin period, H is the Heaviside step function and $g(\varphi)$ is the intrinsic total intensity profile shape. For a power-law spectrum of density inhomogeneities in the ionized ISM τ is expected to have a power-law dependence on frequency ν as

$$\tau(\nu) = \tau_{\nu_\circ} \big(\frac{\nu}{\nu_\circ}\big)^{\alpha}, \tag{6.2}$$

with reference frequency $\nu_\circ$. In all cases, the scattering timescale (133 ms at 2 GHz; see below) dominates the smearing from the process of incoherent dedispersion ($\sim$0.7 ms $(\nu/2\ \mathrm{GHz})^{-3}$), the smearing from an incorrect DM when averaging channels ($\sim$25 μs $(\delta\mathrm{DM/cm^{-3}\ pc})$ $(\nu/2\ \mathrm{GHz})^{-3}$), and the temporal resolution (1.8 ms for 2048 profile bins), so we have not included those modifications of the pulse profile shape in the model. However, deviation from the simple timing models discussed in §6.4.3 (e.g., see Figure 6.4) during any of these epochs could add profile smearing in the integrated profiles (at the level of $\sim$tens of ms — a significant fraction of the scattering timescale). We avoided this source of bias by iterating over the timing model to remove the timing residual on a per-epoch basis.

We model g with a frequency-dependent Gaussian function,

$$g(\nu, \varphi) = A(\nu)\ \exp\big(- 4\ln(2) \frac{(\varphi - \varphi_g(\nu))^2}{\sigma(\nu)^2} \big), \tag{6.3}$$

which is parameterized by its location $\varphi_g(\nu)$, full-width-at-half-maximum (FWHM) $\sigma(\nu)$, and amplitude $A(\nu)$.

As described in Pennucci et al. (2014), each of these parameters nominally has an additional parameter describing its frequency dependence. However, because this combined band has a fractional bandwidth of "only" $\sim$0.5, we assume $\varphi_g(\nu) = \varphi_\circ$ is a frequency-independent value. That is, we assume there is no drift intrinsic to the one component across the band.

Furthermore, when allowing for a frequency-dependent σ, we found no significant evolution, and so we chose also to fix the evolution $\sigma(\nu) = \sigma_\circ$ to be a frequency-independent fit parameter in our final portrait models. This choice was further justified by performing independent per-channel profile fits of a single, scattered Gaussian component and examining the frequency evolution of σ. Also, there are X-band observations for three of these epochs, and in these cases the FWHM of the X-band profiles (all of "Type 3"), when fitted with a single, unscattered Gaussian component, was always within the scatter of those measured from the lower frequency observations. These results are consistent with the weak (or lack of) frequency dependence of σ found in Spitler et al. (2014).

We normalized the intensities of the data to be fit by the maximum profile value in each frequency channel to remove the unusual spectral shape (see §6.4.4). This allowed the Gaussian amplitude to be easily modeled by a power-law function for $A(\nu)$. In all cases, the reduced χ^2 of the fit was <1.1, and a second Gaussian component was never justified by the residuals.

The combination of the quality of the X-band data, the variability of the profile, and the expected value of τ at 8.9 GHz ($\lesssim$1 ms, comparable with our native time resolution) was such that we did not attempt to incorporate this high frequency data into our wideband profile model, nor did we measure the scattering timescale in either the average profile or the single pulses. We refer the reader to Bower et al. (2014) and Spitler et al. (2014) for high frequency scattering measurements of J1745−2900.

Pulse Width & Scattering Parameters

The results from our wideband models are shown in Figure 6.8. There was no significant change in the measured FWHM of the unscattered profile, and our average (frequency-independent) value of 91.9(4) ms = 0.0244(1) rot is also consistent with what is reported in Spitler et al. (2014). The scattering timescale at 2 GHz, $\tau_{2\mathrm{GHz}}$, appears to increase by ∼10% over the four weeks, and the scattering index α deviates from its average value, first to a Kolmogorov value near −4.4 (the dash-dotted line), and then to a much shallower value near −3.0. Both of these results are somewhat peculiar, but similar variations are also reported in Spitler et al. (2014), though they do not discuss the temporal evolution of either quantity. That is, their published values of τ from a variety of epochs and frequencies cannot be unified by a single scattering timescale and index. In fact, their measurements of τ show more scatter over the course of their observations than those presented here, which have some overlap. When the authors combine all of their measurements, they find an average value for α of −3.8(2).

We checked our measurements in two ways. First, we performed conventional profile fits of a single, scattered Gaussian component to each individual frequency channel, independent of any evolutionary constraint. The values of σ, τ, and α for

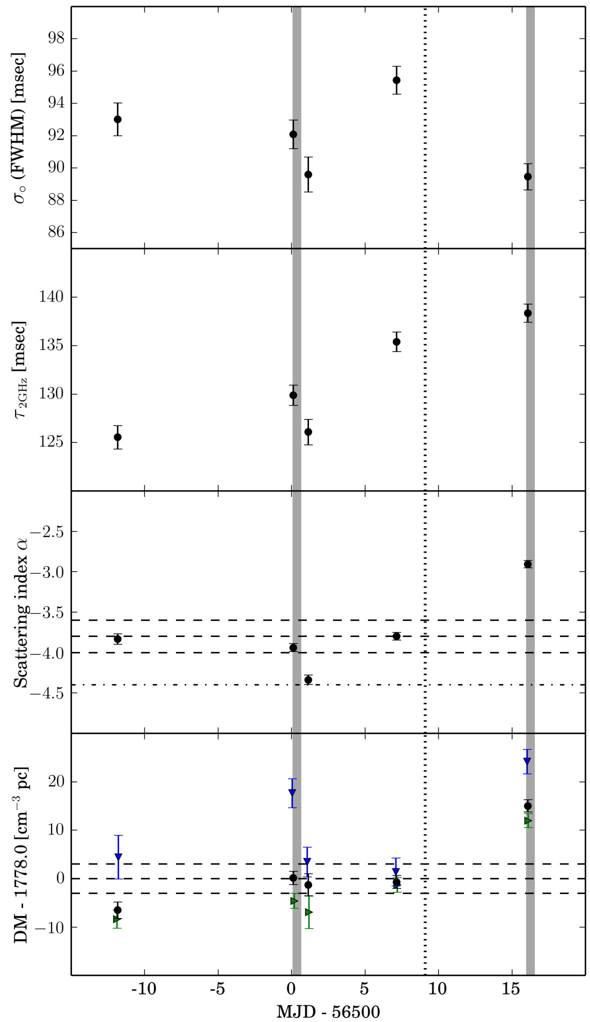

Fig. 6.8 – Our wideband measurements from five epochs. The vertical demarcations are the same as in Figure 6.4. The FWHM showed neither frequency nor temporal dependence. The trend in the scattering timescale τ is less scattered and more precise than the measurements presented in Spitler et al. (2014). The dash-dotted line in the panel for the scattering index α marks the fiducial Kolmogorov value of -4.4, and the dashed lines mark the Spitler et al. (2014) measurement of $-3.8(2)$. The additional markers in the bottom panel are the same as in Figure 6.6: the blue/down-pointing triangles are L-band measurements, and the green/right-pointing triangles are S-band measurements — the dots are their weighted average. The dashed lines here are the Eatough et al. (2013b) DM of 1778(3) cm^{-3} pc. See the text for a discussion of the DM measurements.

each epoch were consistent with what we found by applying the wideband modeling method. In Figure 6.9, we show the measurements of τ measured in this way for the brightest observed epoch (MJD 56516) and over-plot the fitted power-law, which has the most extreme α value of the five epochs. There was nothing unusual about the data from this epoch in terms of RFI, data removal, calibration, or baseline variations. Second, as a check for our average values, we summed all of the data portraits together by coherently stacking the observations (having fit for a phase and DM in each epoch), and fit a single wideband model to the averaged data (with the same constraints as earlier). Using this method, we obtained similar average values: $\sigma = 0.0246(1)$ rot, $\tau_{2\mathrm{GHz}} = 133.0(5)$ ms, and $\alpha = -3.71(2)$, the latter of which is in concert with the average α value from Spitler et al. (2014). Our extrapolated value of $\tau_{1\mathrm{GHz}} = 1.74(3)$ s is only slightly at odds with their average value of $\tau_{1\mathrm{GHz}} = 1.3(2)$ s, which is probably due to the temporal variability of τ. As others have noted (e.g., Bower et al. (2014)), the anticipated value for $\tau_{1\mathrm{GHz}}$ along this line of sight based on empirical relationships, for a DM of 1778 cm^{-3} pc, is about 600$\times$ larger than what is observed (Krishnakumar et al. 2015; Lewandowski et al. 2015a).

Bower et al. (2014) determined $\tau_{8.7\mathrm{GHz}} \lesssim 2$ ms from interferometric measurements of J1745$-$2900's single pulses, implying that a value for α as shallow as -3 is not unbelievable. Furthermore, scattering measurements from two high-DM pulsars discovered near the Galactic Center (both within 0.3° and having DMs 1100–1200 cm^{-3} pc) implied $\alpha = -3.0(3)$ (Johnston et al. 2006). It is not uncommon for pulsars to have $\alpha > -4$, particularly along special lines-of-sight, and it is empirically suggested that the highest DM pulsars may have an average scattering index significantly shallower than -4 (Löhmer et al. 2001, 2004b; Lewandowski et al. 2015a). Note that observing $\alpha \neq -4.4$ does not necessarily imply a non-Kolmogorov spectrum of density inhomogeneities; rather, it could be that a non-thin-screen geometry may be responsible (Cordes & Lazio 2001; Lewandowski et al. 2013).

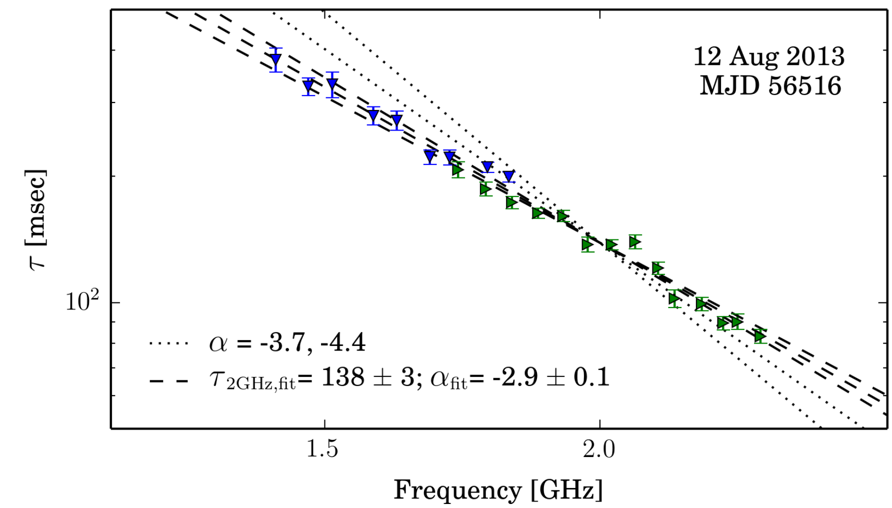

Fig. 6.9 – Independent per-channel measurements of the scattering timescale τ and the fitted scattering index α for the brightest set of L- and S-band observations, on MJD 56516; similar plots from the other four days have a significantly more negative slope. The dashed lines represent our measurement, whereas the dotted lines show our average value of α from wideband modeling and the fiducial Kolmogorov value of -4.4, all with the same value of $\tau_{2\mathrm{GHz}}$. Here, the measurement uncertainties have been inflated by the reduced $\chi^2 \sim 2$.

Dispersion Measures

The bottom panel in Figure 6.8 shows the best-fit DMs as determined by the wideband models in the essentially simultaneous L-band and S-band observations (blue/down-pointing and green/right-pointing triangles, respectively). The black points are the weighted average of the two measurements; there is obviously some variance about the nominal value of 1778(3) cm^{-3} pc, and our overall average value is $\sim$1781(1) cm^{-3} pc. Without exception, the measured L-band DMs are greater than those measured in S-band. There is also only one epoch where the 1σ uncertainties have any overlap; the RMS variance of the differences is $\sim$6 cm^{-3} pc. The absolute DM differences cause residual dispersion on the order of $\lesssim$20 ms $\sim$ 10 bins (for 2048-bin profiles) across the corresponding band, and so they present significant profile deviations. To make sense of the discrepant DMs between the two frequency bands, we consider a number of possibilities.

First, the time-averaged data for each epoch showed few systematics with negligible baseline variations, so we do not believe that data quality was an issue here.

Next, as is well known, the measured absolute DM will be affected by the choice of profile alignment[18]. We can rule out any simple, constant profile evolution as the source of the differing DMs because such a modification introduces a *constant* difference in the DMs; the changes in the measured DM should be the same independent of the choice of alignment. Even if our assumption that there is no intrinsic drift in the location of the (unscattered) profile component across the band is wrong, allowing for a drifting component will still reproduce discrepant DMs; we have verified this by allowing for frequency evolution in the location parameter of the Gaussian component.

A second confounding element from our modeling could be the use of different models for each epoch; if they are all systematically wrong in their alignments or representation, they could be wrong differently. One way to check this is to simply use one fixed model to remake the DM measurements. We used the average portrait model discussed earlier and confirmed that the DMs remain similarly extreme, within $\sim$2 cm^{-3} pc, comparable with the measurement uncertainties. In fact, we tried a large number of fixed and variable portrait models, but never obtained either consistent DMs or DMs with a near constant offset. So, to the extent that τ and/or α are measurably changing, we are justified in keeping them as free parameters for each epoch's model.

Similarly, the known profile variability that is seen in all of the radio magnetars

[18]For example, DMs are significantly biased when either assuming a constant profile shape in the presence of scattering, or aligning scattered profiles by conventional methods because the convolution of the ISM pulse broadening function with a profile of finite width introduces a delay that is a function of the scattering timescale (i.e. frequency). This is partly why proper wideband modeling is necessary.

could also play a role when using either a fixed or variable portrait model. However, besides the flux density, any underlying profile shape changes either with time or frequency are masked by the large level of scattering. As mentioned, the FWHM does not seem to change significantly in either time or frequency. Furthermore, the three X-band observations taken during these epochs show no large profile changes, and are all of the "Type 3" shape.

Next we can ask whether or not the slight asynchronicity could have any effect; that is, could the DM change so significantly on $\sim$hour timescales? We will return to this question below, but it is not an uncommon *a priori* assumption to expect that the observed DM does not change between observations separated by $\lesssim 4$ hr.

One could also ask if the method by which we measure the DMs introduces a systematic error, where the error may depend on the exact values of τ and α, or even the spectral shape. To answer this, we performed a number of Monte Carlo simulations. In the simulations, we used the models from MJDs 56500 and 56516, which have the most extreme values for the difference between the DMs, and the most extreme values for τ and α, respectively. For each trial, we made fake L- and S-band observations by appropriately constructing the model for that band and scaling each frequency channel's amplitude to match the spectral shape. We then added random frequency-dependent white noise to the model at the same level as measured from the data portraits and finally dispersed the fake data with a DM of 1778 cm^{-3} pc. Visual inspection verified that the fake data were faithfully rendered. We used the same method to measure the DM (and phase), which is described in detail in Pennucci et al. (2014). In summary, the measured DMs were always in accord and unbiased, and the uncertainties were accurately estimated. We conclude that the measurement method produces accurate DMs, independent of the model parameters, provided that the model for the data is accurate.

We assume in our measurements that the phase offsets ($\Delta\phi$) incurred by finite-frequency signals due to propagation through the ionized interstellar medium scale as predicted by the usual cold-plasma dispersion law such that $\Delta\phi \propto \frac{\mathrm{DM}}{P_s}\,\nu^{-2}$. This is certainly the case to first-order even over large, low frequency bandwidths (Hassall et al. 2012). However, to the extent that we understand the ISM to be *in*homogeneous — after all, we do observe pulse broadening — then it is anticipated that the simple ν^{-2} dependence will be an insufficient description at some level for broadband DM measurements. When an inhomogeneous medium causes multi-path propagation of radio waves where the path depends on frequency, the sampled column density of free electrons (the DM) will also be a function of frequency. Thus, we are left with the intriguing possible explanation that the DM inconsistencies we are seeing are the consequence of imposing a ν^{-2} dispersion law onto a frequency-dependent DM (DM(ν)) due to an inhomogeneous ISM[19].

To our knowledge, the most recent claim for having observed frequency-dependent DMs was reported in Ahuja et al. (2007) for the slow, low DM pulsars B0329+54 and B1642−03, although they observed lower DMs at lower frequencies. However, the authors only made one set of simultaneous pairs of dual-frequency measurements per pulsar. We argue that to confidently segregate the effects of profile evolution, DM variations with time (DM(t)), DM(ν), and other potential confounding factors, many epochs of simultaneous, wideband (large fractional bandwidth) observations of a stable, preferably high DM pulsar need to be made. A similar recommendation was recently made by Cordes et al. (2015) in their detailed study of frequency-dependent DMs, which makes theoretical predictions for the characteristic timescales and sizes

[19]This is opposed to other supposed origins of DM(ν) relating to magnetospheric propagation effects or magnetic sweepback, which would likely have different statistics from an ISM induced DM(ν); see Hassall et al. (2012) or Ahuja et al. (2007) for an overview.

of DM(ν) effects.

In their treatment of the problem, Cordes et al. (2015) predict the minimum scale of DM variations about a mean value,

$$\overline{\mathrm{DM}}_{\mathrm{rms}} \sim \phi_F^2/\lambda r_e \sim 3.84 \times 10^{-8}\ \mathrm{cm}^{-3}\ \mathrm{pc}\ \nu_{\mathrm{GHz}}\ \phi_F^2, \tag{6.4}$$

where ν_{GHz} is the frequency in GHz and ϕ_F is the size of the phase perturbations over the Fresnel scale, $l_F = \sqrt{(cD)/(2\pi\nu)}$, for the speed of light c and source distance D. For J1745−2900, which is in the strong scattering regime, ϕ_F will be very large. We estimate it from their prescription,

$$\phi_F(\nu) \approx 9.6\ \mathrm{rad}\ \left(\frac{\nu/\Delta\nu_d}{100}\right)^{5/12}, \tag{6.5}$$

where $\Delta\nu_d$ is the scintillation bandwidth, which is readily estimated from our scattering measurements as $\sim 1.16/(2\pi\tau(\nu))$. For 1.4 and 2.4 GHz, we find $\overline{\mathrm{DM}}_{\mathrm{rms}} \sim$ 10 and 5 cm^{-3} pc, respectively. These can be compared to the RMS DM values as measured in L- and S-band of $\sim$ 9 and 7 cm^{-3} pc, respectively. The characteristic spatial size for the DM differences near 2 GHz will be several Fresnel scales, which can be converted to a characteristic time by using the recently measured proper motion of 236 km s^{-1} (Bower et al. 2015). For our range of frequencies, the characteristic timescale associated with the Fresnel scale size is ∼3 hr, comparable to the separation between the observations on a given epoch. Therefore, it may be that the small temporal gap between the observations contributes somewhat to the difference in the DMs, but we certainly do expect that the DMs vary significantly on different days, separated by many Fresnel timescales.

Finally, Cordes et al. (2015) make a prediction for the observed RMS difference between DMs at frequencies ν and ν',

$$\sigma_{\overline{\mathrm{DM}}}(\nu,\nu') \approx 4.42 \times 10^{-5}\ \mathrm{cm}^{-3}\ \mathrm{pc}\ F_\beta(r)\left(\frac{\nu\phi_F^2}{1000}\right), \tag{6.6}$$

where we have ignored a geometric factor of order unity and the function F_β contains the frequency dependence for $r \equiv \nu/\nu'$, given the power-law index β for the wavenumber spectrum of density inhomogeneities. For ν = 2.4 GHz and ν' = 1.4 GHz, $\sigma_{\overline{\mathrm{DM}}} \sim 4$ cm^{-3} pc, compared to our observed RMS difference of ∼6 cm^{-3} pc.

That the predicted and observed values are similar may be coincidence, but we note the corroborating facts that J1745−2900 is the highest DM pulsar, is relatively bright, highly scattered, has a simple, easily modeled profile, and does not show significant profile evolution or stochastic profile variability (at least in these observations). Furthermore, we verified that our measurement method produces inconsistent (and biased) DMs between the bands by introducing non-ν^{-2} phase delays into our fake data simulations described earlier. After ruling out the other potential sources for the inconsistent DMs, we suggest that J1745−2900 may have an observable frequency-dependent dispersion measure.

A potential counter argument is that over many Fresnel timescales, one expects the sign of the DM differences to change, such that the observed low frequency DM becomes smaller than the high frequency DM. Between the small number and low density of epochs, the potentially incorrect portrait model, and ISM uncertainties (the predictions here are based on a thin-screen model with a Kolmogorov spectrum of density perturbations, which is partly supported by the findings in Bower et al. (2014)), it is conceivable that this observation is not inconsistent with a frequency-dependent DM as described.

Determining whether or not a difference in DM as seen in two frequency bands is intrinsically a DM(ν) effect is complicated by the issues described above, and with only five measurements we obviously cannot draw any definite or statistical conclusions, but future studies could potentially disentangle the evolution of DM(t,ν), the profile, and other ISM parameters. One strategy, as Cordes et al. (2015) note, is to model the frequency dependence of the dispersive delays as something other than ν^{-2}. This should be done for many epochs, at least as long as the timescale for refractive scintillation, over which time the specific frequency dependence of the average DM remains stable. For J1745−2900, this timescale is potentially many years.

6.5 Summary & Discussion

In this paper we have presented multi-epoch, multi-frequency wideband GBT observations of the Galactic Center radio magnetar J1745−2900 at 1.5, 2.0, and 8.9 GHz from the first ∼100 days after it was discovered. After its initial X-ray burst on 25 April 2013, J1745−2900 underwent two additional bursts in the course of our observations. For two epochs, during which time we collected data from three radio bands, we also have simultaneous X-ray observations taken with *Chandra*. An analysis of the radio data, as well as a joint analysis with the X-ray data, yielded a few noteworthy results.

1. We found no anomalous radio bursts or giant-pulse-like individual pulses in any of the X-band observations. Similarly, the smooth transitioning of the X-band profile between three broad categories seems to have also been unperturbed by the *Swift*-detected X-ray bursts.

2. Our simple radio timing analysis corroborates the findings of Kaspi et al. (2014), which are also supported by Coti Zelati et al. (2015). We presented the absolute

alignment of the three radio and 0.3-8 keV profiles. The near-alignment of the radio components with the X-ray profile is similar to the two other radio magnetars that have published alignments. We also make note of a possible transient X-ray feature from Coti Zelati et al. (2015) because of its proximity to the phase of radio emission located ~ 0.2 in phase preceding the peak in the X-ray profile.

3. The evolution of our early radio flux measurements, showing a relatively stable growth from around the time of the initial outburst, is consistent with the continued GBT X-band observations presented in Lynch et al. (2014) and with what they have called a "stable state"[20]. The combination of the gradual flux evolution with the simple timing and profile variability results leads us to extrapolate J1745−2900's "stable state" back to the time of its initial burst.

4. The shape of J1745−2900's low frequency radio spectrum is potentially positive or flat, whereas it shows a "typical" spectral index of ~ -1.4 between $\sim$2 and 9 GHz, at least during a brief period $\sim$100 days after its initial outburst, around the times of two later X-ray bursts. This steep spectral index might indicate a different magnetospheric configuration during these times, although the evolving spectra may be a result of environmental factors and free-free absorption (Lewandowski et al. 2015b). The possible variability of γ means that dedicated observations covering several higher frequency bands need to be carried out over many epochs to confirm this (cf. Torne et al. 2015).

5. We made wideband models of J1745−2900's low frequency radio "portrait" to measure the scattering timescale, scattering index, and the DM as a function of

[20]While our observations were taken over a shorter range of time (about a third), our cadence of observations is comparable to theirs taken during the "erratic state", the onset of which was apparently unrelated to X-ray bursts.

time. Our average measurements are consistent with what has been published in Spitler et al. (2014), though the ISM parameters may be variable. Time-variable scattering parameters would complicate the predicted sensitivities of future pulsar surveys of the Galactic Center. Lastly, we make a suggestion that our discrepant, nearly simultaneously determined DMs are a manifestation of an ISM-induced frequency-dependent dispersion measure, and that future observations could make a case study out of J1745−2900 to investigate DM(ν) — provided the pulsar remains visible and stable.

J1745−2900 shares several (but not all) properties with the other three radio magnetars, J1809−1943, J1550−5418, and J1622−4950 (Camilo et al. 2006, 2007b; Levin et al. 2010). Common properties of the pulsed radio emission from magnetars are: a) a delay in the appearance of the radio emission after the X-ray outburst onset, b) variable pulse profiles and radio flux on timescales from hours to days, c) a large rotational (spin-down) luminosity with respect to the quiescent X-ray luminosity, d) a decrease of the radio flux as the X-ray flux decays, and e) a flat radio spectrum over a wide range of frequencies. J1745−2900 grossly shares the first three properties with the rest of its class. However, while in all other cases the radio flux was observed to decay as the X-ray outburst was fading, the long-term radio and X-ray flux evolution of J1745−2900 is at variance with this trend. The radio flux shows a re-brightening hundreds of days after the outburst onset and the X-ray emission is decaying very slowly, challenging current crustal cooling models Coti Zelati et al. (2015). Furthermore, the recently published flux measurements by Torne et al. (2015) taken one year after those presented here suggest that the 8.35 GHz flux remained stable at the $\sim$3 mJy level over thirty days. Another interesting peculiarity of this radio magnetar was the steep (and possibly free-free absorbed) radio spectrum seen in our observations, though the more recent observations in Torne et al. (2015) suggest that the spectrum has since flattened.

Of particular interest is J1745−2900's low quiescent luminosity compared to its high rotational power ($L_{\rm X,qui}/L_{\rm rot} < 1$; Rea et al. 2013). This peculiarity of the four radio magnetars, which is at variance with canonical magnetars (for which the fact that $L_{\rm X,qui}/L_{\rm rot} > 1$ has always been used as proof of their magnetically dominated emission (Mereghetti & Stella 1995; Thompson & Duncan 1995; Mereghetti 2008)), has been viewed as evidence for a similar mechanism powering the radio emission from magnetars and normal pulsars alike. In fact, while normal radio pulsars have primarily dipolar-dominated magnetic fields ($B_{\rm p}$), magnetars have a substantial toroidal component (B_ϕ) that is present in both the internal and external fields. This toroidal component is the main reason for their quiescent X-ray luminosities, hot surface temperatures, flaring emission, and outburst activity (Thompson et al. 2002; Beloborodov 2009). For a fixed dipolar field, the internal toroidal field has no significant effect on the luminosity unless $B_\phi > B_{\rm p}$, as is the case for most magnetars (Viganò et al. 2013). Both radio magnetars and high-B radio pulsars have systematically lower toroidal fields and higher rotational energies than typical magnetars; this is in agreement with the former being fainter in quiescence and having a softer X-ray spectrum (a lower crustal toroidal field results in less heating produced by Joule dissipation in the crust, Pons et al. 2009). As for the energy powering the radio emission, simulations of high dipolar field pulsars that have a small toroidal component showed that the particle acceleration and subsequent ignition of the cascade process could proceed as it does in normal pulsars, successfully reaching the open-field line region and generating pulsed radio emission (Medin & Lai 2010). On the other hand, for an extremely strong toroidal component, it is expected that the particle cascades cannot reach the open-field lines due to the powerful currents formed as a consequence of the twisted magnetosphere. Radio magnetars might lie in between, having a high enough rotational energy to power pair cascades as in normal pulsars, but also having toroidal components lower than typical magnetars, resulting in lower quiescent X-ray luminosities.

In the above picture, the possible radio flux increase, the steep spectrum, and the slow cooling of the X-ray outburst might be explained by the presence of a strongly twisted bundle, which can account for the radio emission and the additional heating by particles slamming onto the surface. If the radio emission is generated by acceleration of particles only in this part of the magnetosphere, then the radio flux and the X-ray flux might be unrelated. In particular, untwisting of the bundle during the outburst decay might induce fewer currents blocking the pair cascade generation, hence more radio emission from this region. However, these are only speculative, plausible hypotheses. Proof of this scenario would need a longer monitoring of the radio and X-ray emission, as well as detailed magnetohydrodynamical simulations of particle acceleration and pair cascades in a strongly magnetized and twisted bundle.

Chapter 7

Summary

“Pulsars are cool. Seriously.” —Scott M. Ransom

7.1 Summary & Future Studies

In this dissertation, we have demonstrated the utility of simple two-dimensional phase-frequency model portraits to make time-of-arrival (TOA) and dispersion measure (DM) measurements in folded profile pulsar data. Pulse profile evolution obviously does not need to be modeled with Gaussian components — indeed, we have seen some of the shortcomings of this approach — but modeling profile evolution is unavoidable in the era of truly ultimate, wideband pulsar instruments. There are many possible alternatives to Gaussian components, and others have already explored much more sophisticated, Bayesian approaches to estimating profile evolution simultaneously with DM variations and the timing model (Lentati et al. 2015a).

One long term goal for the portrait modeling is to eventually make use of the full polarization information we have from the NANOGrav data set. A study investigating the components in the polarization profiles of the PPTA pulsars was recently published, which may have paved the way for the other PTAs to follow suit (Dai et al. 2015). Ultimately, if the model parameters can be informative for studies of the pulsar magnetosphere, then that is what we should strive for.

The two cases in which we found evidence of something resembling a chromatic DM, J1643−1224 and J1745−2900, certainly warrant further investigation. In any event, the implication is that with a sufficiently good portrait model, a bright enough pulsar, a wide enough band, and the flexibility to look for non-ν^{-2} dispersion, theoretical ISM effects may soon be able to be probed observationally. A grand demonstration would be one in which DM(t,ν) is completely segregated from profile evolution and perhaps other ISM effects. Our immediate next goal is to develop a streamlined way to fit for any combination of the TOA, DM, scattering timescale, and scattering index on a per-epoch basis. However, our methods as they stand should prove useful in a variety of systems that have, e.g., variable DM behavior on short timescales. One such pulsar is the globular cluster binary MSP Ter5A, which shows irregular eclipses and DM behavior (A. Bilous, private communication).

Additionally, we made some unintended contributions in our study of the Galactic Center magnetar. Namely, our observations corroborate the models presented in Lewandowski et al. (2015b) and imply that we have seen some free-free absorbing material either in the ejecta from J1745−2900 or in its environment. We are preparing to work with the authors to share data and scrutinize this matter further. It is unfortunate that additional observations would likely be uninformative, unless the magnetar has another X-ray burst.

Our preliminary results from our Shapiro delay investigations are currently being followed up. Even just two additional pulsar masses with precisions of 0.1–0.2 $M_\odot$ would constitute an $\sim 8\%$ increase in the number of pulsar masses known.

Finally, our first wholesale implementation of the wideband protocol to the NANOGrav 9-year profile data set proved instructive. Indeed, it highlighted the shortcomings of our Gaussian models and suggested that we take a different approach to modeling the noise in our data. However, we were pleased to find that most of the red noise parameters, as well as the overall timing residuals and dispersion measures, were all in agreement. We believe we have at least paved the way for these wideband methods to wholly supplant the channelized TOA approaches of the past. Our next critical test, once we have obtained better noise parameters, is to determine gravitational wave limits from our data set and compare them with others'.

With the advent of new technologies and telescopes (e.g., CHIME, FAST, meerKAT), and the growing sizes and cooperation of PTA collaborations, the next decade of pulsar astronomy promises to hold long-awaited answers and unexpected surprises.

References

Ahuja, A. L., Gupta, Y., Mitra, D., & Kembhavi, A. K. 2005, MNRAS, 357, 1013

Ahuja, A. L., Mitra, D., & Gupta, Y. 2007, MNRAS, 377, 677

Anders, E. & Grevesse, N. 1989, Geochim. Cosmochim. Acta, 53, 197

Anderson, G. E., Gaensler, B. M., Slane, P. O., Rea, N., Kaplan, D. L., Posselt, B., Levin, L., Johnston, S., Murray, S. S., Brogan, C. L., Bailes, M., Bates, S., Benjamin, R. A., Bhat, N. D. R., Burgay, M., Burke-Spolaor, S., Chakrabarty, D., D'Amico, N., Drake, J. J., Esposito, P., Grindlay, J. E., Hong, J., Israel, G. L., Keith, M. J., Kramer, M., Lazio, T. J. W., Lee, J. C., Mauerhan, J. C., Milia, S., Possenti, A., Stappers, B., & Steeghs, D. T. H. 2012, ApJ, 751, 53

Archibald, A. M., Stairs, I. H., Ransom, S. M., Kaspi, V. M., Kondratiev, V. I., Lorimer, D. R., McLaughlin, M. A., Boyles, J., Hessels, J. W. T., Lynch, R., van Leeuwen, J., Roberts, M. S. E., Jenet, F., Champion, D. J., Rosen, R., Barlow, B. N., Dunlap, B. H., & Remillard, R. A. 2009, Science, 324, 1411

Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., Chamberlin, S., Chatterjee, S., Christy, B., Cordes, J. M., Cornish, N., Crowter, K., Demorest, P. B., Dolch, T., Ellis, J. A., Ferdman, R. D., Fonseca, E., Garver-Daniels, N., Gonzalez, M. E., Jenet, F. A., Jones, G., Jones, M., Kaspi, V. M., Koop, M., Lazio, T. J. W., Lam,

M. T., Levin, L., Lommen, A. N., Lorimer, D. R., Luo, J., Lynch, R. S., Madison, D., McLaughlin, M. A., McWilliams, S. T., Nice, D. J., Palliyaguru, N., Pennucci, T. T., Ransom, S. M., Siemens, X., Stairs, I. H., Stinebring, D. R., Stovall, K., Swiggum, J. K., Vallisneri, M., van Haasteren, R., Wang, Y., & Zhu, W. 2015a, ArXiv e-prints

Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., Chamberlin, S. J., Chatterjee, S., Christy, B., Cordes, J. M., Cornish, N. J., Demorest, P. B., Deng, X., Dolch, T., Ellis, J. A., Ferdman, R. D., Fonseca, E., Garver-Daniels, N., Jenet, F., Jones, G., Kaspi, V. M., Koop, M., Lam, M. T., Lazio, T. J. W., Levin, L., Lommen, A. N., Lorimer, D. R., Luo, J., Lynch, R. S., Madison, D. R., McLaughlin, M. A., McWilliams, S. T., Nice, D. J., Palliyaguru, N., Pennucci, T. T., Ransom, S. M., Siemens, X., Stairs, I. H., Stinebring, D. R., Stovall, K., Swiggum, J., Vallisneri, M., van Haasteren, R., Wang, Y., & Zhu, W. W. 2015b, ArXiv e-prints

Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., Chamberlin, S. J., Chatterjee, S., Cordes, J. M., Demorest, P. B., Deng, X., Dolch, T., Ellis, J. A., Ferdman, R. D., Garver-Daniels, N., Jenet, F., Jones, G., Kaspi, V. M., Koop, M., Lam, M. T., Lazio, T. J. W., Lommen, A. N., Lorimer, D. R., Luo, J., Lynch, R. S., Madison, D. R., McLaughlin, M. A., McWilliams, S. T., Nice, D. J., Palliyaguru, N., Pennucci, T. T., Ransom, S. M., Sesana, A., Siemens, X., Stairs, I. H., Stinebring, D. R., Stovall, K., Swiggum, J., Vallisneri, M., van Haasteren, R., Wang, Y., Zhu, W. W., & NANOGrav Collaboration. 2014, ApJ, 794, 141

Baade, W. & Zwicky, F. 1934, Physical Review, 46, 76

Backer, D. C., Hama, S., van Hook, S., & Foster, R. S. 1993, ApJ, 404, 636

Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., & Goss, W. M. 1982, Nature, 300, 615

Balucinska-Church, M. & McCammon, D. 1992, ApJ, 400, 699

Bates, S. D., Johnston, S., Lorimer, D. R., Kramer, M., Possenti, A., Burgay, M., Stappers, B., Keith, M. J., Lyne, A., Bailes, M., McLaughlin, M. A., O'Brien, J. T., & Hobbs, G. 2011, MNRAS, 411, 1575

Bates, S. D., Lorimer, D. R., & Verbiest, J. P. W. 2013, MNRAS, 431, 1352

Beloborodov, A. M. 2009, ApJ, 703, 1044

Bertotti, B., Carr, B. J., & Rees, M. J. 1983, MNRAS, 203, 945

Bhat, N. D. R., Cordes, J. M., Camilo, F., Nice, D. J., & Lorimer, D. R. 2004, ApJ, 605, 759

Bilous, A. V., Pennucci, T. T., Demorest, P., & Ransom, S. M. 2015, ApJ, 803, 83

Bower, G. C., Deller, A., Demorest, P., Brunthaler, A., Eatough, R., Falcke, H., Kramer, M., Lee, K. J., & Spitler, L. 2014, ApJ, 780, L2

Bower, G. C., Deller, A., Demorest, P., Brunthaler, A., Falcke, H., Moscibrodzka, M., O'Leary, R. M., Eatough, R. P., Kramer, M., Lee, K. J., Spitler, L., Desvignes, G., Rushton, A. P., Doeleman, S., & Reid, M. J. 2015, ApJ, 798, 120

Boyles, J., Lynch, R. S., Ransom, S. M., Stairs, I. H., Lorimer, D. R., McLaughlin, M. A., Hessels, J. W. T., Kaspi, V. M., Kondratiev, V. I., Archibald, A., Berndsen, A., Cardoso, R. F., Cherry, A., Epstein, C. R., Karako-Argaman, C., McPhee, C. A., Pennucci, T., Roberts, M. S. E., Stovall, K., & van Leeuwen, J. 2013, ApJ, 763, 80

Bracewell, R. N. 2000, The Fourier transform and its applications

Callanan, P. J., Garnavich, P. M., & Koester, D. 1998, MNRAS, 298, 207

Camilo, F., Cognard, I., Ransom, S. M., Halpern, J. P., Reynolds, J., Zimmerman, N., Gotthelf, E. V., Helfand, D. J., Demorest, P., Theureau, G., & Backer, D. C. 2007a, ApJ, 663, 497

Camilo, F., Ransom, S. M., Halpern, J. P., & Reynolds, J. 2007b, ApJ, 666, L93

Camilo, F., Ransom, S. M., Halpern, J. P., Reynolds, J., Helfand, D. J., Zimmerman, N., & Sarkissian, J. 2006, Nature, 442, 892

Camilo, F., Ransom, S. M., Peñalver, J., Karastergiou, A., van Kerkwijk, M. H., Durant, M., Halpern, J. P., Reynolds, J., Thum, C., Helfand, D. J., Zimmerman, N., & Cognard, I. 2007c, ApJ, 669, 561

Camilo, F., Reynolds, J., Johnston, S., Halpern, J. P., & Ransom, S. M. 2008, ApJ, 679, 681

Chen, J. L. & Wang, H. G. 2014, ApJS, 215, 11

Chennamangalam, J. & Lorimer, D. R. 2014, MNRAS, 440, L86

Cognard, I., Guillemot, L., Johnson, T. J., Smith, D. A., Venter, C., Harding, A. K., Wolff, M. T., Cheung, C. C., Donato, D., Abdo, A. A., Ballet, J., Camilo, F., Desvignes, G., Dumora, D., Ferrara, E. C., Freire, P. C. C., Grove, J. E., Johnston, S., Keith, M., Kramer, M., Lyne, A. G., Michelson, P. F., Parent, D., Ransom, S. M., Ray, P. S., Romani, R. W., Saz Parkinson, P. M., Stappers, B. W., Theureau, G., Thompson, D. J., Weltevrede, P., & Wood, K. S. 2011, ApJ, 732, 47

Cognard, I. & Lestrade, J.-F. 1997, A&A, 323, 211

Comella, J. M., Craft, H. D., Lovelace, R. V. E., & Sutton, J. M. 1969, Nature, 221, 453

Cordes, J. M. 1978, ApJ, 222, 1006

—. in prep.

Cordes, J. M. & Lazio, T. J. W. 1997, ApJ, 475, 557

—. 2001, ApJ, 549, 997

Cordes, J. M. & Shannon, R. M. 2010, ArXiv e-prints

Cordes, J. M., Shannon, R. M., & Stinebring, D. R. 2015, ArXiv e-prints

Coti Zelati, F., Rea, N., Papitto, A., Viganò, D., Pons, J. A., Turolla, R., Esposito, P., Haggard, D., Baganoff, F. K., Ponti, G., Israel, G. L., Campana, S., Torres, D. F., Tiengo, A., Mereghetti, S., Perna, R., Zane, S., Mignani, R. P., Possenti, A., & Stella, L. 2015, MNRAS, 449, 2685

Craft, Jr., H. D. 1970, PhD thesis, Cornell University.

Dai, S., Hobbs, G., Manchester, R. N., Kerr, M., Shannon, R. M., van Straten, W., Mata, A., Bailes, M., Bhat, N. D. R., Burke-Spolaor, S., Coles, W. A., Johnston, S., Keith, M. J., Levin, Y., Osłowski, S., Reardon, D., Ravi, V., Sarkissian, J. M., Tiburzi, C., Toomey, L., Wang, H. G., Wang, J.-B., Wen, L., Xu, R. X., Yan, W. M., & Zhu, X.-J. 2015, MNRAS, 449, 3223

Davis, J. E. 2001, ApJ, 562, 575

Degenaar, N., Reynolds, M. T., Miller, J. M., Kennea, J. A., & Wijnands, R. 2013, The Astronomer's Telegram, 5006, 1

Dembska, M., Basu, R., Kijak, J., & Lewandowski, W. 2015, MNRAS, 449, 1869

Dembska, M., Kijak, J., Jessner, A., Lewandowski, W., Bhattacharyya, B., & Gupta, Y. 2014, MNRAS, 445, 3105

Demorest, P. B. 2007, PhD thesis, University of California, Berkeley

Demorest, P. B., Ferdman, R. D., Gonzalez, M. E., Nice, D., Ransom, S., Stairs, I. H., Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., Chamberlin, S. J., Cordes, J. M., Ellis, J., Finn, L. S., Freire, P., Giampanis, S., Jenet, F., Kaspi, V. M., Lazio, J., Lommen, A. N., McLaughlin, M., Palliyaguru, N., Perrodin, D., Shannon, R. M., Siemens, X., Stinebring, D., Swiggum, J., & Zhu, W. W. 2013, ApJ, 762, 94

Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, Nature, 467, 1081

Deneva, J. S., Cordes, J. M., & Lazio, T. J. W. 2009, ApJ, 702, L177

Detweiler, S. 1979, ApJ, 234, 1100

Dexter, J. & O'Leary, R. M. 2014, ApJ, 783, L7

DuPlain, R., Ransom, S., Demorest, P., Brandt, P., Ford, J., & Shelton, A. L. 2008, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7019, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series

Dyks, J. & Rudak, B. 2015, MNRAS, 446, 2505

Eatough, R., Karuppusamy, R., Kramer, M., Klein, B., Champion, D., Kraus, A., Keane, E., Bassa, C., Lyne, A., Lazarus, P., Verbiest, J., Freire, P., Brunthaler, A., & Falcke, H. 2013a, The Astronomer's Telegram, 5040, 1

Eatough, R. P., Falcke, H., Karuppusamy, R., Lee, K. J., Champion, D. J., Keane, E. F., Desvignes, G., Schnitzeler, D. H. F. M., Spitler, L. G., Kramer, M., Klein, B., Bassa, C., Bower, G. C., Brunthaler, A., Cognard, I., Deller, A. T., Demorest, P. B., Freire, P. C. C., Kraus, A., Lyne, A. G., Noutsos, A., Stappers, B., & Wex, N. 2013b, Nature, 501, 391

Ellis, J. A. 2013, Classical and Quantum Gravity, 30, 224004

—. 2014, PhD thesis, University of Wisconsin, Milwaukee

Ellis, J. A., Siemens, X., & van Haasteren, R. 2013, ApJ, 769, 63

Foster, R. S. & Backer, D. C. 1990, ApJ, 361, 300

Foster, R. S., Fairhead, L., & Backer, D. C. 1991, ApJ, 378, 687

Freeman, P., Doe, S., & Siemiginowska, A. 2001, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4477, Astronomical Data Analysis, ed. J.-L. Starck & F. D. Murtagh, 76–87

Freire, P. C. C. & Wex, N. 2010, MNRAS, 409, 199

Fruscione, A., McDowell, J. C., Allen, G. E., Brickhouse, N. S., Burke, D. J., Davis, J. E., Durham, N., Elvis, M., Galle, E. C., Harris, D. E., Huenemoerder, D. P., Houck, J. C., Ishibashi, B., Karovska, M., Nicastro, F., Noble, M. S., Nowak, M. A., Primini, F. A., Siemiginowska, A., Smith, R. K., & Wise, M. 2006, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 6270, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series

Garmire, G. P., Bautz, M. W., Ford, P. G., Nousek, J. A., & Ricker, Jr., G. R. 2003, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4851, X-Ray and Gamma-Ray Telescopes and Instruments for Astronomy., ed. J. E. Truemper & H. D. Tananbaum, 28–44

Genzel, R., Eisenhauer, F., & Gillessen, S. 2010, Reviews of Modern Physics, 82, 3121

Gold, T. 1968, Nature, 218, 731

Goldreich, P. & Julian, W. H. 1969, ApJ, 157, 869

Gotthelf, E. V., Petre, R., & Vasisht, G. 1999, ApJ, 514, L107

Guillemot, L., Freire, P. C. C., Cognard, I., Johnson, T. J., Takahashi, Y., Kataoka, J., Desvignes, G., Camilo, F., Ferrara, E. C., Harding, A. K., Janssen, G. H., Keith, M., Kerr, M., Kramer, M., Parent, D., Ransom, S. M., Ray, P. S., Saz Parkinson, P. M., Smith, D. A., Stappers, B. W., & Theureau, G. 2012, MNRAS, 422, 1294

Halpern, J. P., Gotthelf, E. V., Reynolds, J., Ransom, S. M., & Camilo, F. 2008, ApJ, 676, 1178

Hankins, T. H. & Rickett, B. J. 1975, Methods in Computational Physics, 14, 55

Hassall, T. E., Stappers, B. W., Hessels, J. W. T., Kramer, M., Alexov, A., Anderson, K., Coenen, T., Karastergiou, A., Keane, E. F., Kondratiev, V. I., Lazaridis, K., van Leeuwen, J., Noutsos, A., Serylak, M., Sobey, C., Verbiest, J. P. W., Weltevrede, P., Zagkouris, K., Fender, R., Wijers, R. A. M. J., Bähren, L., Bell, M. E., Broderick, J. W., Corbel, S., Daw, E. J., Dhillon, V. S., Eislöffel, J., Falcke, H., Grießmeier, J.-M., Jonker, P., Law, C., Markoff, S., Miller-Jones, J. C. A., Osten, R., Rol, E., Scaife, A. M. M., Scheers, B., Schellart, P., Spreeuw, H., Swinbank, J., ter Veen, S., Wise, M. W., Wijnands, R., Wucknitz, O., Zarka, P., Asgekar, A., Bell, M. R., Bentum, M. J., Bernardi, G., Best, P., Bonafede, A., Boonstra, A. J., Brentjens, M., Brouw, W. N., Brüggen, M., Butcher, H. R., Ciardi, B., Garrett, M. A., Gerbers, M., Gunst, A. W., van Haarlem, M. P., Heald, G., Hoeft, M., Holties, H., de Jong, A., Koopmans, L. V. E., Kuniyoshi, M., Kuper, G., Loose, G. M., Maat, P., Masters, J., McKean, J. P., Meulman, H., Mevius, M., Munk, H., Noordam, J. E., Orrú, E., Paas, H., Pandey-Pommier, M., Pandey, V. N., Pizzo, R., Polatidis, A., Reich, W., Röttgering, H., Sluman, J., Steinmetz, M., Sterks, C. G. M., Tagger, M., Tang, Y., Tasse, C., Vermeulen, R., van Weeren, R. J., Wijnholds, S. J., & Yatawatta, S. 2012, A&A, 543, A66

Hellings, R. W. & Downs, G. S. 1983, ApJ, 265, L39

Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, Nature, 217, 709

Hobbs, G. 2013, Classical and Quantum Gravity, 30, 224007

Hotan, A. W., Bailes, M., & Ord, S. M. 2005, MNRAS, 362, 1267

—. 2006, MNRAS, 369, 1502

Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, Proc. Astron. Soc. Aust., 21, 302

Jacoby, B. A., Bailes, M., van Kerkwijk, M. H., Ord, S., Hotan, A., Kulkarni, S. R., & Anderson, S. B. 2003, ApJ, 599, L99

Jaffe, A. H. & Backer, D. C. 2003, ApJ, 583, 616

Johnson, T. J., Guillemot, L., Kerr, M., Cognard, I., Ray, P. S., Wolff, M. T., Bégin, S., Janssen, G. H., Romani, R. W., Venter, C., Grove, J. E., Freire, P. C. C., Wood, M., Cheung, C. C., Casandjian, J. M., Stairs, I. H., Camilo, F., Espinoza, C. M., Ferrara, E. C., Harding, A. K., Johnston, S., Kramer, M., Lyne, A. G., Michelson, P. F., Ransom, S. M., Shannon, R., Smith, D. A., Stappers, B. W., Theureau, G., & Thorsett, S. E. 2013, ApJ, 778, 106

Johnston, S., Kramer, M., Lorimer, D. R., Lyne, A. G., McLaughlin, M., Klein, B., & Manchester, R. N. 2006, MNRAS, 373, L6

Kaspi, V. M., Archibald, R. F., Bhalerao, V., Dufour, F., Gotthelf, E. V., An, H., Bachetti, M., Beloborodov, A. M., Boggs, S. E., Christensen, F. E., Craig, W. W., Grefenstette, B. W., Hailey, C. J., Harrison, F. A., Kennea, J. A., Kouveliotou, C., Madsen, K. K., Mori, K., Markwardt, C. B., Stern, D., Vogel, J. K., & Zhang, W. W. 2014, ApJ, 786, 84

Kaspi, V. M., Taylor, J. H., & Ryba, M. F. 1994, ApJ, 428, 713

Keith, M. J., Coles, W., Shannon, R. M., Hobbs, G. B., Manchester, R. N., Bailes, M., Bhat, N. D. R., Burke-Spolaor, S., Champion, D. J., Chaudhary, A., Hotan, A. W., Khoo, J., Kocz, J., Osłowski, S., Ravi, V., Reynolds, J. E., Sarkissian, J., van Straten, W., & Yardley, D. R. B. 2013, MNRAS, 429, 2161

Keith, M. J., Johnston, S., Levin, L., & Bailes, M. 2011, MNRAS, 416, 346

Kennea, J. A., Burrows, D. N., Cummings, J., Kouveliotou, C., Degenaar, N., Reynolds, M. T., Miller, J. M., & Wijnands, R. 2013a, The Astronomer's Telegram, 5124, 1

Kennea, J. A., Burrows, D. N., Cummings, J., Krimm, H. A., Barthelmy, S., Kouveliotou, C., Degenaar, N., Reynolds, M. T., Miller, J. M., & Wijnands, R. 2013b, The Astronomer's Telegram, 5254, 1

Kennea, J. A., Krimm, H., Barthelmy, S., Gehrels, N., Markwardt, C., Cummings, J., Marshall, F., Sakamoto, T., Degenaar, N., Reynolds, M. T., Miller, J. M., & Kouveliotou, C. 2013c, The Astronomer's Telegram, 5009, 1

Kijak, J., Lewandowski, W., Maron, O., Gupta, Y., & Jessner, A. 2011, A&A, 531, A16

Kijak, J., Tarczewski, L., Lewandowski, W., & Melikidze, G. 2013, ApJ, 772, 29

Komesaroff, M. M. 1970, Nature, 225, 612

Kramer, M. & Champion, D. J. 2013, Classical and Quantum Gravity, 30, 224009

Kramer, M., Lange, C., Lorimer, D. R., Backer, D. C., Xilouris, K. M., Jessner, A., & Wielebinski, R. 1999, ApJ, 526, 957

Kramer, M., Xilouris, K. M., Lorimer, D. R., Doroshenko, O., Jessner, A., Wielebinski, R., Wolszczan, A., & Camilo, F. 1998, ApJ, 501, 270

Krishnakumar, M. A., Mitra, D., Naidu, A., Joshi, B. C., & Manoharan, P. K. 2015, ApJ, 804, 23

Lam, M. T., Cordes, J. M., Chatterjee, S., & Dolch, T. 2015, ApJ, 801, 130

Lange, C., Camilo, F., Wex, N., Kramer, M., Backer, D. C., Lyne, A. G., & Doroshenko, O. 2001, MNRAS, 326, 274

Large, M. I., Vaughan, A. E., & Mills, B. Y. 1968, Nature, 220, 340

Lattimer, J. M. 2012, Annual Review of Nuclear and Particle Science, 62, 485

Lattimer, J. M. & Prakash, M. 2005, Physical Review Letters, 94, 111101

Lazaridis, K., Jessner, A., Kramer, M., Stappers, B. W., Lyne, A. G., Jordan, C. A., Serylak, M., & Zensus, J. A. 2008, MNRAS, 390, 839

Lazaridis, K., Wex, N., Jessner, A., Kramer, M., Stappers, B. W., Janssen, G. H., Desvignes, G., Purver, M. B., Cognard, I., Theureau, G., Lyne, A. G., Jordan, C. A., & Zensus, J. A. 2009, MNRAS, 400, 805

Lazarus, P., Brazier, A., Hessels, J. W. T., Karako-Argaman, C., Kaspi, V. M., Lynch, R., Madsen, E., Patel, C., Ransom, S. M., Scholz, P., Swiggum, J., Zhu, W. W., Allen, B., Bogdanov, S., Camilo, F., Cardoso, F., Chatterjee, S., Cordes, J. M., Crawford, F., Deneva, J. S., Ferdman, R., Freire, P. C. C., Jenet, F. A., Knispel, B., Lee, K. J., van Leeuwen, J., Lorimer, D. R., Lyne, A. G., McLaughlin, M. A., Siemens, X., Spitler, L. G., Stairs, I. H., Stovall, K., & Venkataraman, A. 2015, ArXiv e-prints

Lee, K. J., Bassa, C. G., Janssen, G. H., Karuppusamy, R., Kramer, M., Liu, K., Perrodin, D., Smits, R., Stappers, B. W., van Haasteren, R., & Lentati, L. 2014, MNRAS, 441, 2831

Lentati, L., Alexander, P., & Hobson, M. P. 2013, ArXiv e-prints

—. 2015a, MNRAS, 447, 2159

Lentati, L., Taylor, S. R., Mingarelli, C. M. F., Sesana, A., Sanidas, S. A., Vecchio, A., Caballero, R. N., Lee, K. J., van Haasteren, R., Babak, S., Bassa, C. G., Brem, P., Burgay, M., Champion, D. J., Cognard, I., Desvignes, G., Gair, J. R., Guillemot, L., Hessels, J. W. T., Janssen, G. H., Karuppusamy, R., Kramer, M., Lassus, A., Lazarus, P., Liu, K., Osłowski, S., Perrodin, D., Petiteau, A., Possenti, A., Purver, M. B., Rosado, P. A., Smits, R., Stappers, B., Theureau, G., Tiburzi, C., & Verbiest, J. P. W. 2015b, ArXiv e-prints

Levin, L., Bailes, M., Bates, S., Bhat, N. D. R., Burgay, M., Burke-Spolaor, S., D'Amico, N., Johnston, S., Keith, M., Kramer, M., Milia, S., Possenti, A., Rea, N., Stappers, B., & van Straten, W. 2010, ApJ, 721, L33

Levin, L., Bailes, M., Bates, S. D., Bhat, N. D. R., Burgay, M., Burke-Spolaor, S., D'Amico, N., Johnston, S., Keith, M. J., Kramer, M., Milia, S., Possenti, A., Stappers, B., & van Straten, W. 2012, MNRAS, 422, 2489

Levin, L. e. a. in prep., ApJ

Lewandowski, W., Dembska, M., Kijak, J., & Kowalińska, M. 2013, MNRAS, 434, 69

Lewandowski, W., Kowalińska, M., & Kijak, J. 2015a, MNRAS, 449, 1570

Lewandowski, W., Rożko, K., Kijak, J., & Melikidze, G. I. 2015b, ArXiv e-prints

Liu, K., Desvignes, G., Cognard, I., Stappers, B. W., Verbiest, J. P. W., Lee, K. J., Champion, D. J., Kramer, M., Freire, P. C. C., & Karuppusamy, R. 2014, MNRAS, 443, 3752

Löhmer, O., Kramer, M., Driebe, T., Jessner, A., Mitra, D., & Lyne, A. G. 2004a, A&A, 426, 631

Löhmer, O., Kramer, M., Mitra, D., Lorimer, D. R., & Lyne, A. G. 2001, ApJ, 562, L157

Löhmer, O., Mitra, D., Gupta, Y., Kramer, M., & Ahuja, A. 2004b, A&A, 425, 569

Lommen, A. N. 2001, PhD thesis, University of California, Berkeley

Lommen, A. N. & Demorest, P. 2013, Classical and Quantum Gravity, 30, 224001

Lorimer, D. R. & Kramer, M. 2005, Handbook of Pulsar Astronomy, ed. R. Ellis, J. Huchra, S. Kahn, G. Rieke, & P. B. Stetson (The Press Syndicate of the University of Cambridge)

Lynch, R. S., Archibald, R. F., Kaspi, V. M., & Scholz, P. 2014, ArXiv e-prints

Lynch, R. S., Boyles, J., Ransom, S. M., Stairs, I. H., Lorimer, D. R., McLaughlin, M. A., Hessels, J. W. T., Kaspi, V. M., Kondratiev, V. I., Archibald, A. M., Berndsen, A., Cardoso, R. F., Cherry, A., Epstein, C. R., Karako-Argaman, C., McPhee, C. A., Pennucci, T., Roberts, M. S. E., Stovall, K., & van Leeuwen, J. 2013, ApJ, 763, 81

Lyne, A. & Graham-Smith, F. 2012, Pulsar Astronomy

Lyne, A. G., Brinklow, A., Middleditch, J., Kulkarni, S. R., & Backer, D. C. 1987, Nature, 328, 399

Macquart, J.-P. & Kanekar, N. 2015, ArXiv e-prints

Macquart, J.-P., Kanekar, N., Frail, D. A., & Ransom, S. M. 2010, ApJ, 715, 939

Maitia, V., Lestrade, J.-F., & Cognard, I. 2003, ApJ, 582, 972

Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, AJ, 129, 1993

McLaughlin, M. A. 2013, Classical and Quantum Gravity, 30, 224008

—. 2014, ArXiv e-prints

Medin, Z. & Lai, D. 2010, MNRAS, 406, 1379

Mereghetti, S. 2008, A&A Rev., 15, 225

—. 2013, Brazilian Journal of Physics, 43, 356

Mereghetti, S., Pons, J., & Melatos, A. 2015, ArXiv e-prints

Mereghetti, S. & Stella, L. 1995, ApJ, 442, L17

Mori, K., Gotthelf, E. V., Barriere, N. M., Hailey, C. J., Harrison, F. A., Kaspi, V. M., Tomsick, J. A., & Zhang, S. 2013a, The Astronomer's Telegram, 5020, 1

Mori, K., Gotthelf, E. V., Zhang, S., An, H., Baganoff, F. K., Barrière, N. M., Beloborodov, A. M., Boggs, S. E., Christensen, F. E., Craig, W. W., Dufour, F., Grefenstette, B. W., Hailey, C. J., Harrison, F. A., Hong, J., Kaspi, V. M., Kennea, J. A., Madsen, K. K., Markwardt, C. B., Nynka, M., Stern, D., Tomsick, J. A., & Zhang, W. W. 2013b, ApJ, 770, L23

Nord, M. E., Lazio, T. J. W., Kassim, N. E., Goss, W. M., & Duric, N. 2004, ApJ, 601, L51

Olausen, S. A. & Kaspi, V. M. 2014, ApJS, 212, 6

Özel, F., Psaltis, D., Narayan, R., & Santos Villarreal, A. 2012, ApJ, 757, 55

Pacini, F. 1967, Nature, 216, 567

—. 1968, Nature, 219, 145

Papitto, A., Ferrigno, C., Bozzo, E., Rea, N., Pavan, L., Burderi, L., Burgay, M., Campana, S., di Salvo, T., Falanga, M., Filipović, M. D., Freire, P. C. C., Hessels, J. W. T., Possenti, A., Ransom, S. M., Riggio, A., Romano, P., Sarkissian, J. M., Stairs, I. H., Stella, L., Torres, D. F., Wieringa, M. H., & Wong, G. F. 2013, Nature, 501, 517

Pennucci, T. T., Demorest, P. B., & Ransom, S. M. 2014, ApJ, 790, 93

Pennucci, T. T., Possenti, A., Esposito, P., Rea, N., Haggard, D., Baganoff, F. K., Burgay, M., Coti Zelati, F., Israel, G. L., & Minter, A. 2015, ArXiv e-prints

Platania, P., Bensadoun, M., Bersanelli, M., de Amici, G., Kogut, A., Levin, S., Maino, D., & Smoot, G. F. 1998, ApJ, 505, 473

Pons, J. A., Miralles, J. A., & Geppert, U. 2009, A&A, 496, 207

Rajagopal, M. & Romani, R. W. 1995, ApJ, 446, 543

Rankin, J. M., Comella, J. M., Craft, Jr., H. D., Richards, D. W., Campbell, D. B., & Counselman, III, C. C. 1970, ApJ, 162, 707

Rankin, J. M., Counselman, III, C. C., & Richards, D. W. 1971, AJ, 76, 686

Ransom, S. M., Ray, P. S., Camilo, F., Roberts, M. S. E., Çelik, Ö., Wolff, M. T., Cheung, C. C., Kerr, M., Pennucci, T., DeCesar, M. E., Cognard, I., Lyne, A. G., Stappers, B. W., Freire, P. C. C., Grove, J. E., Abdo, A. A., Desvignes, G., Donato, D., Ferrara, E. C., Gehrels, N., Guillemot, L., Gwon, C., Harding, A. K., Johnston, S., Keith, M., Kramer, M., Michelson, P. F., Parent, D., Saz Parkinson, P. M., Romani, R. W., Smith, D. A., Theureau, G., Thompson, D. J., Weltevrede, P., Wood, K. S., & Ziegler, M. 2011, ApJ, 727, L16

Ray, P. S., Abdo, A. A., Parent, D., Bhattacharya, D., Bhattacharyya, B., Camilo, F., Cognard, I., Theureau, G., Ferrara, E. C., Harding, A. K., Thompson, D. J., Freire, P. C. C., Guillemot, L., Gupta, Y., Roy, J., Hessels, J. W. T., Johnston, S., Keith, M., Shannon, R., Kerr, M., Michelson, P. F., Romani, R. W., Kramer, M., McLaughlin, M. A., Ransom, S. M., Roberts, M. S. E., Saz Parkinson, P. M., Ziegler, M., Smith, D. A., Stappers, B. W., Weltevrede, P., & Wood, K. S. 2012, ArXiv e-prints

Rea, N., Esposito, P., Pons, J. A., Turolla, R., Torres, D. F., Israel, G. L., Possenti, A., Burgay, M., Viganò, D., Papitto, A., Perna, R., Stella, L., Ponti, G., Baganoff, F. K., Haggard, D., Camero-Arranz, A., Zane, S., Minter, A., Mereghetti, S., Tiengo, A., Schödel, R., Feroci, M., Mignani, R., & Götz, D. 2013, ApJ, 775, L34

Rea, N., Pons, J. A., Torres, D. F., & Turolla, R. 2012, ApJ, 748, L12

Richards, D. W. & Comella, J. M. 1969, Nature, 222, 551

Roberts, J. A. & Richards, D. W. 1971, Nature Physical Science, 231, 25

Roy, J., Ray, P. S., Bhattacharyya, B., Stappers, B., Chengalur, J. N., Deneva, J., Camilo, F., Johnson, T. J., Wolff, M., Hessels, J. W. T., Bassa, C. G., Keane, E. F., Ferrara, E. C., Harding, A. K., & Wood, K. S. 2015, ApJ, 800, L12

Ruderman, M. A. & Sutherland, P. G. 1975, ApJ, 196, 51

Ryba, M. F. & Taylor, J. H. 1991, ApJ, 371, 739

Shannon, R. M. & Cordes, J. M. 2010, ApJ, 725, 1607

Shannon, R. M. & Johnston, S. 2013, MNRAS, 435, L29

Shapiro, I. I. 1964, Physical Review Letters, 13, 789

Siemens, X., Ellis, J., Jenet, F., & Romano, J. D. 2013, Classical and Quantum Gravity, 30, 224015

Siemion, A., Bailes, M., Bower, G., Chennamangalam, J., Cordes, J., Demorest, P., Deneva, J., Desvignes, G., Ford, J., Frail, D., Jones, G., Kramer, M., Lazio, J., Lorimer, D., McLaughlin, M., Ransom, S., Roshi, A., Wagner, M., Werthimer, D., & Wharton, R. 2013, in IAU Symposium, Vol. 291, Proceedings of the International Astronomical Union, ed. J. van Leeuwen, 57–57

Spitler, L. G., Lee, K. J., Eatough, R. P., Kramer, M., Karuppusamy, R., Bassa, C. G., Cognard, I., Desvignes, G., Lyne, A. G., Stappers, B. W., Bower, G. C., Cordes, J. M., Champion, D. J., & Falcke, H. 2014, ApJ, 780, L3

Staelin, D. H. & Reifenstein, III, E. C. 1968, Science, 162, 1481

Stovall, K., Lynch, R. S., Ransom, S. M., Archibald, A. M., Banaszak, S., Biwer, C. M., Boyles, J., Dartez, L. P., Day, D., Ford, A. J., Flanigan, J., Garcia, A., Hessels, J. W. T., Hinojosa, J., Jenet, F. A., Kaplan, D. L., Karako-Argaman, C., Kaspi, V. M., Kondratiev, V. I., Leake, S., Lorimer, D. R., Lunsford, G., Martinez, J. G., Mata, A., McLaughlin, M. A., Roberts, M. S. E., Rohr, M. D., Siemens, X., Stairs, I. H., van Leeuwen, J., Walker, A. N., & Wells, B. L. 2014, ApJ, 791, 67

Sturrock, P. A. 1971, ApJ, 164, 529

Tauris, T. M., Langer, N., & Kramer, M. 2011, MNRAS, 416, 2130

Tauris, T. M. & Savonije, G. J. 1999, A&A, 350, 928

Taylor, J. H. 1992, Royal Society of London Philosophical Transactions Series A, 341, 117

Thompson, C. & Duncan, R. C. 1995, MNRAS, 275, 255

Thompson, C., Lyutikov, M., & Kulkarni, S. R. 2002, ApJ, 574, 332

Thorsett, S. E. 1991, ApJ, 377, 263

Torne, P., Eatough, R. P., Karuppusamy, R., Kramer, M., Paubert, G., Klein, B., Desvignes, G., Champion, D. J., Wiesemeyer, H., Kramer, C., Spitler, L. G., Thum, C., Güsten, R., Schuster, K. F., & Cognard, I. 2015, ArXiv e-prints

Toscano, M., Sandhu, J. S., Bailes, M., Manchester, R. N., Britton, M. C., Kulkarni, S. R., Anderson, S. B., & Stappers, B. W. 1999, MNRAS, 307, 925

van Haasteren, R. & Levin, Y. 2013, MNRAS, 428, 1147

van Haasteren, R. & Vallisneri, M. 2014, Phys. Rev. D, 90, 104012

—. 2015, MNRAS, 446, 1170

van Straten, W. 2013, ApJS, 204, 13

van Straten, W., Demorest, P., & Oslowski, S. 2012, Astronomical Research and Technology, 9, 237

Verbiest, J. P. W., Bailes, M., Coles, W. A., Hobbs, G. B., van Straten, W., Champion, D. J., Jenet, F. A., Manchester, R. N., Bhat, N. D. R., Sarkissian, J. M., Yardley, D., Burke-Spolaor, S., Hotan, A. W., & You, X. P. 2009, MNRAS, 400, 951

Viganò, D., Rea, N., Pons, J. A., Perna, R., Aguilera, D. N., & Miralles, J. A. 2013, MNRAS, 434, 123

Xilouris, K. M., Kramer, M., Jessner, A., Wielebinski, R., & Timofeev, M. 1996, A&A, 309, 481

You, X. P., Hobbs, G., Coles, W. A., Manchester, R. N., Edwards, R., Bailes, M., Sarkissian, J., Verbiest, J. P. W., van Straten, W., Hotan, A., Ord, S., Jenet, F., Bhat, N. D. R., & Teoh, A. 2007, MNRAS, 378, 493

Zhu, W. W., Stairs, I. H., Demorest, P. B., Nice, D. J., Ellis, J. A., Ransom, S. M., Arzoumanian, Z., Crowter, K., Dolch, T., Ferdman, R. D., Fonseca, E., Gonzalez, M. E., Jones, G., Jones, M., Lam, M. T., Levin, L., McLaughlin, M. A., Pennucci, T., Stovall, K., & Swiggum, J. 2015, ArXiv e-prints