Time-Domain Astronomy

Jayanta Roy
NCRA-TIFR

NCRA-IUCAA Graduate School

2nd April 2026

First part of the answer is simple - millisecond pulsars are long-lived, so we should see more of them.

And being in a binary is the only way for the pulsar to become a millisecond pulsar...

Movie credit:
John Rowe
Animations

This is scenatio is called "the recycling of a pulsar".

When the companion star reaches the evolutional stage of the red giant it transfers the matter, through the accretion disc, and the pulsar - even if it was a dead neutron star - starts to spin-up.

The pulsar can reach millisecond periods.

At the same time the orbit is circularized and tightened.

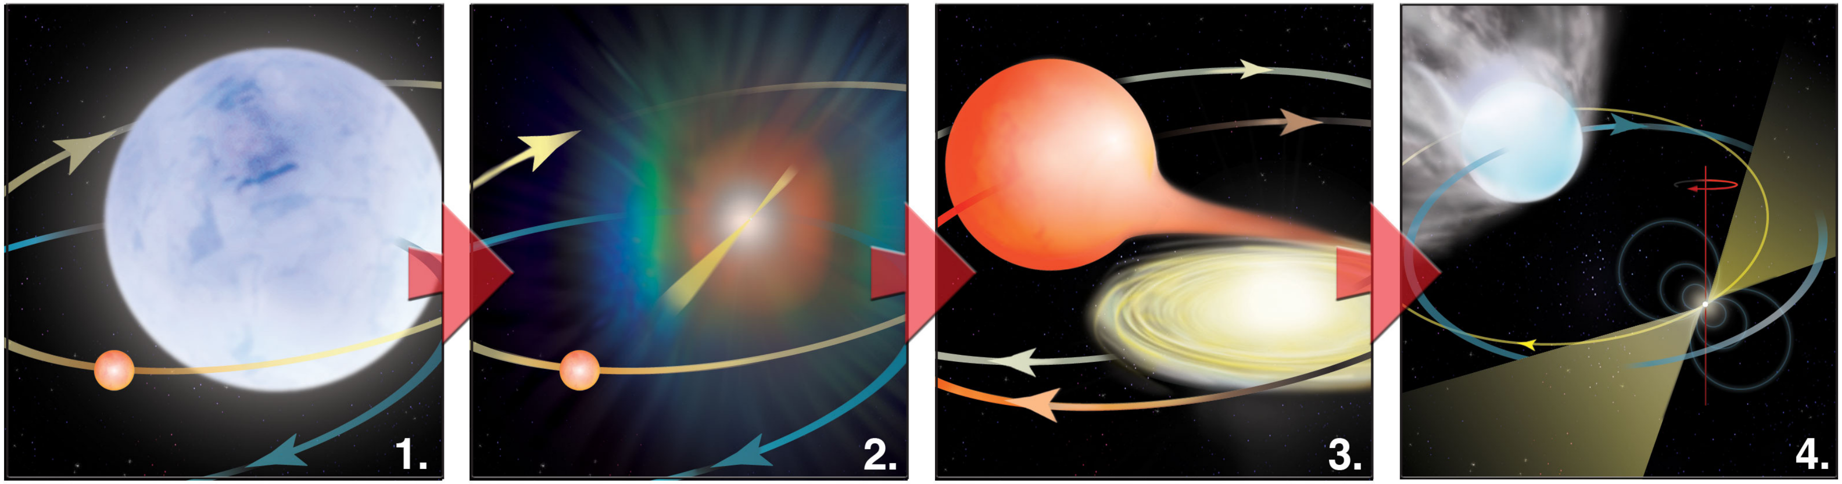

You cannot see the pulsar at this stage - it is completely surrounded by hot plasma, and radio waves cannot pass through it. We see such systems as Low-Mas X-ray Binaries (LMXB).

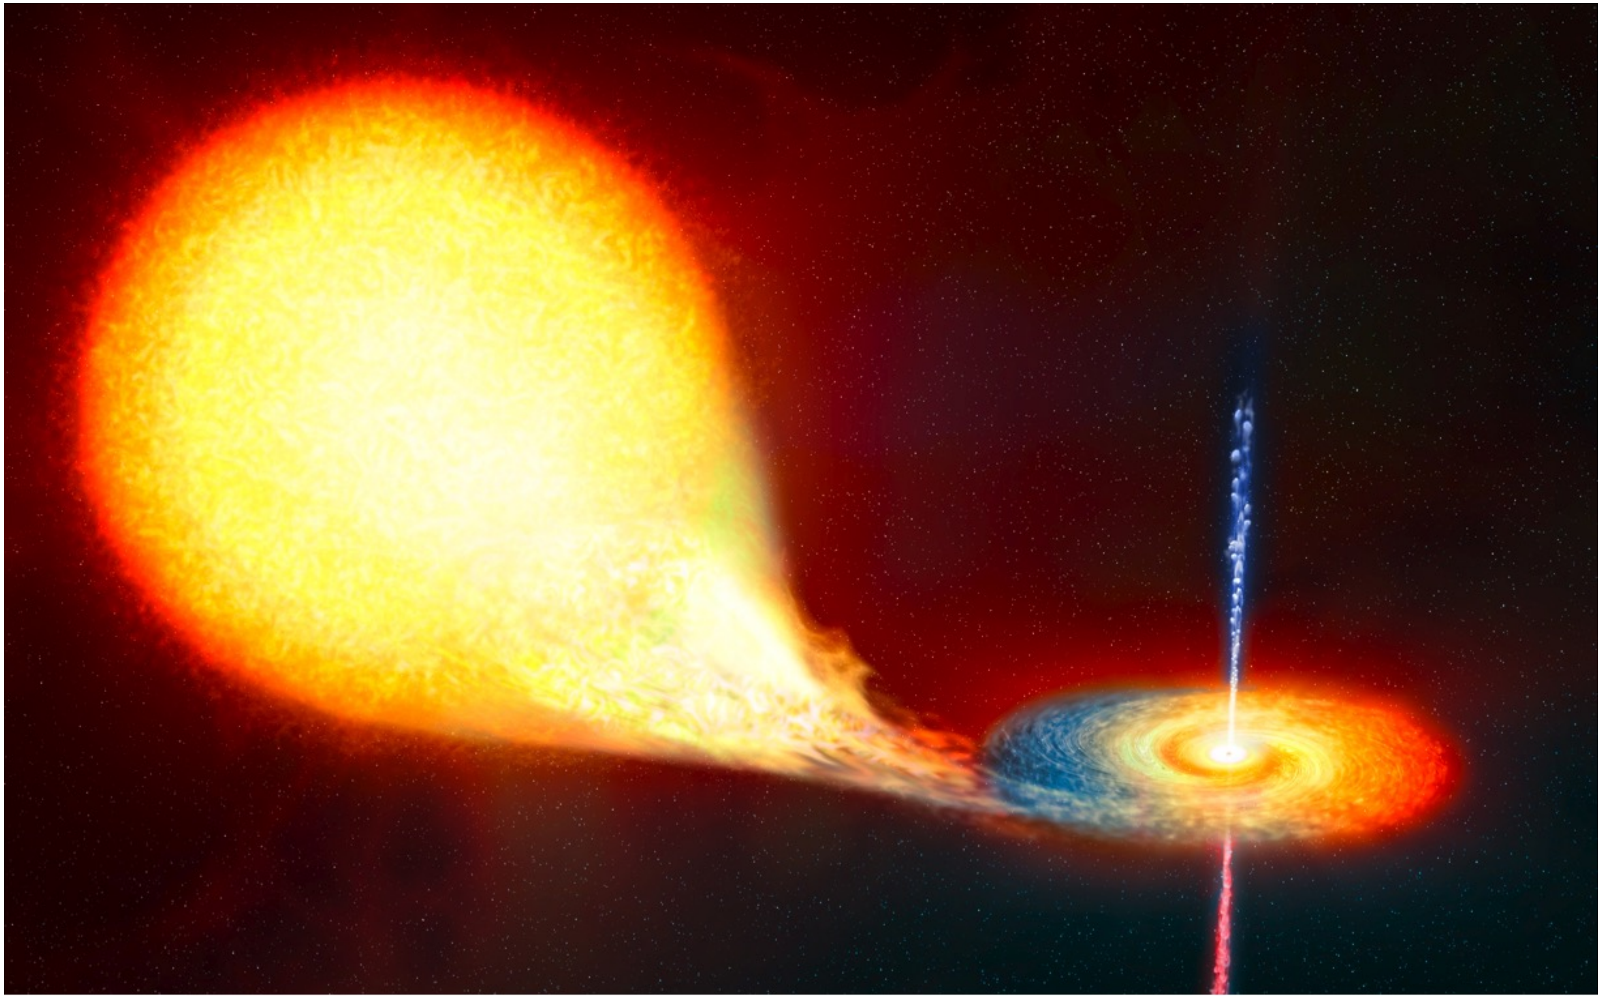

If the pulsar is closer, the radiowaves have shorter way to travel - and they come early. If it is furter away - they come late.

The maximum delay is of course the pulsar semi-major axis divided by the speed of light.

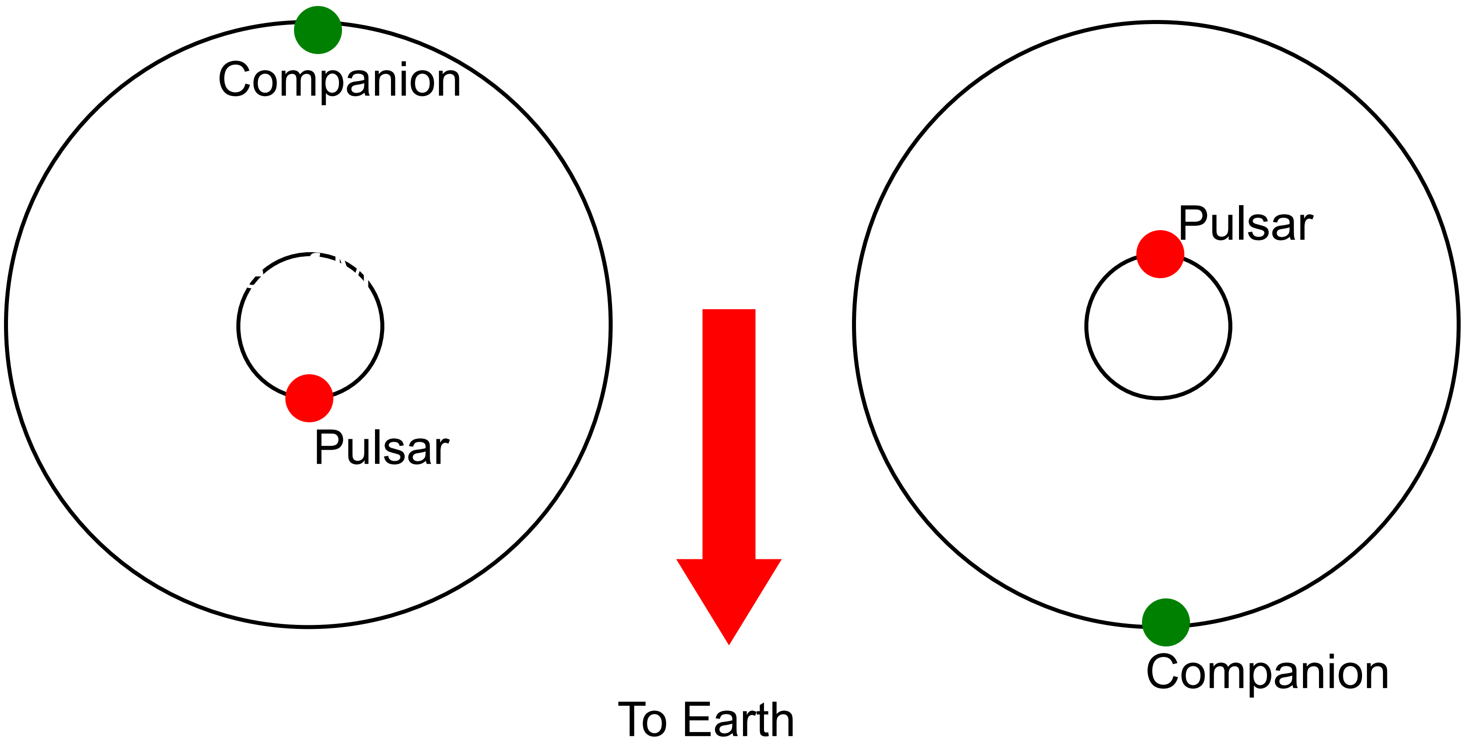

But the timing can tell us more than that. But first, let us recall the basics of celestial mechanics...

There are five basic keplerian parameters describing the pulsar orbit:

a - the semi major axis
P_b - the period of a binary
e - the eccentricity
T_0 - the time of periastron passage
ω - the longitude of periastron.

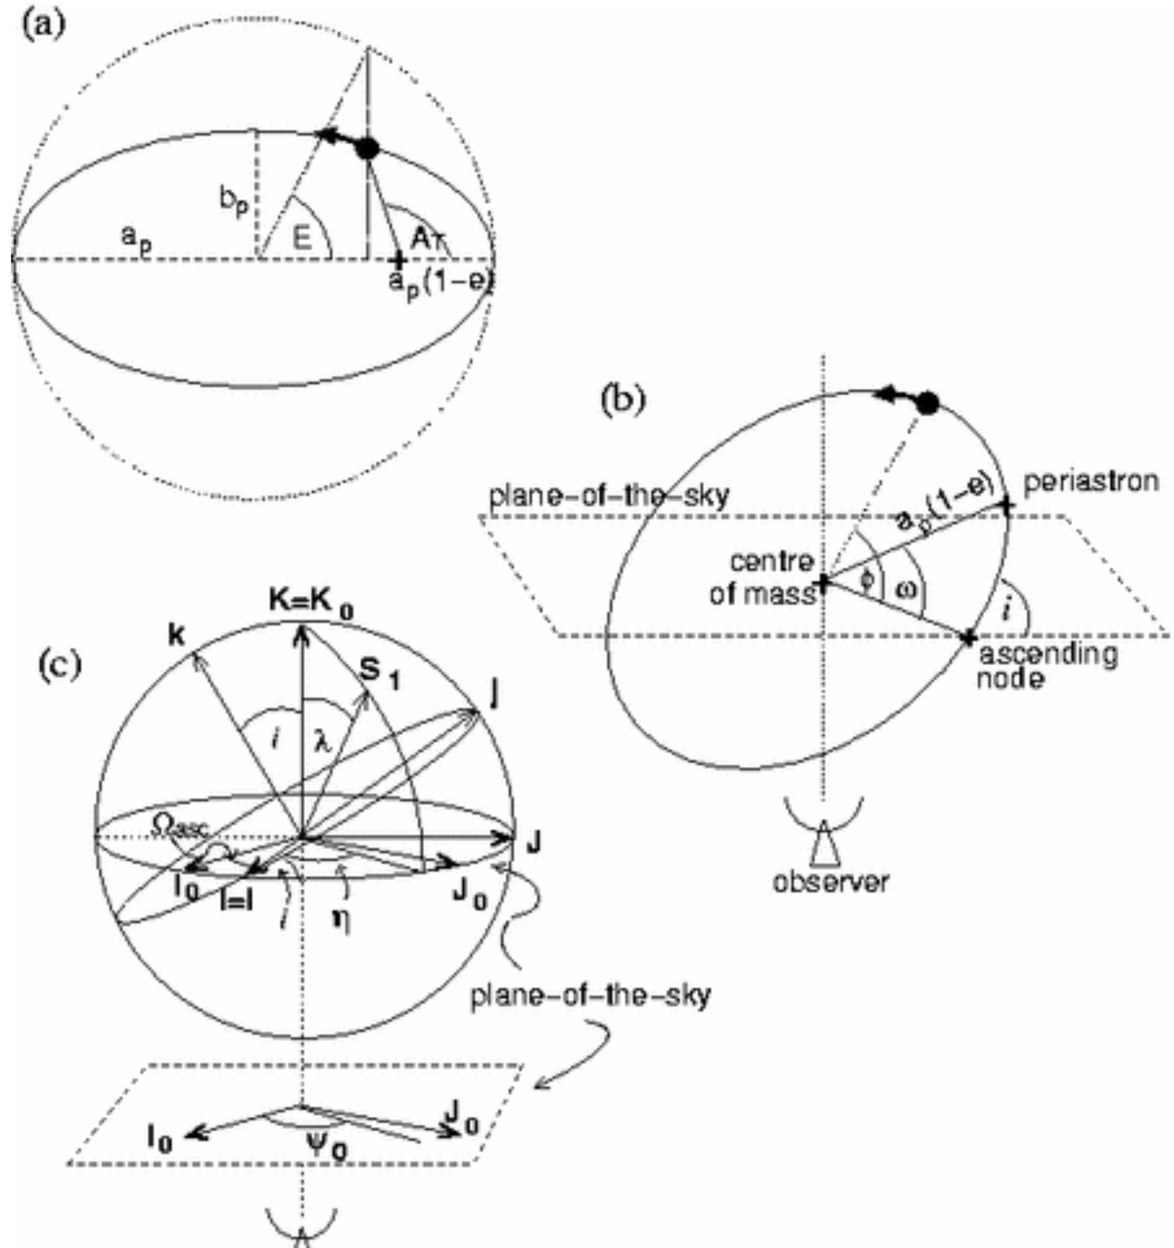

Taken from "Handbook of Pulsar Astronomy" by Lorimer & Kramer

And there is one angle, that we usually cannot tell - the inclination of the orbit *i*.

This results in the fact, that we cannot measure directly the pulsar semi-major axis only its projection to the line of sight. This parameter is usually called *x*, and of course $x = a^* \sin i$.

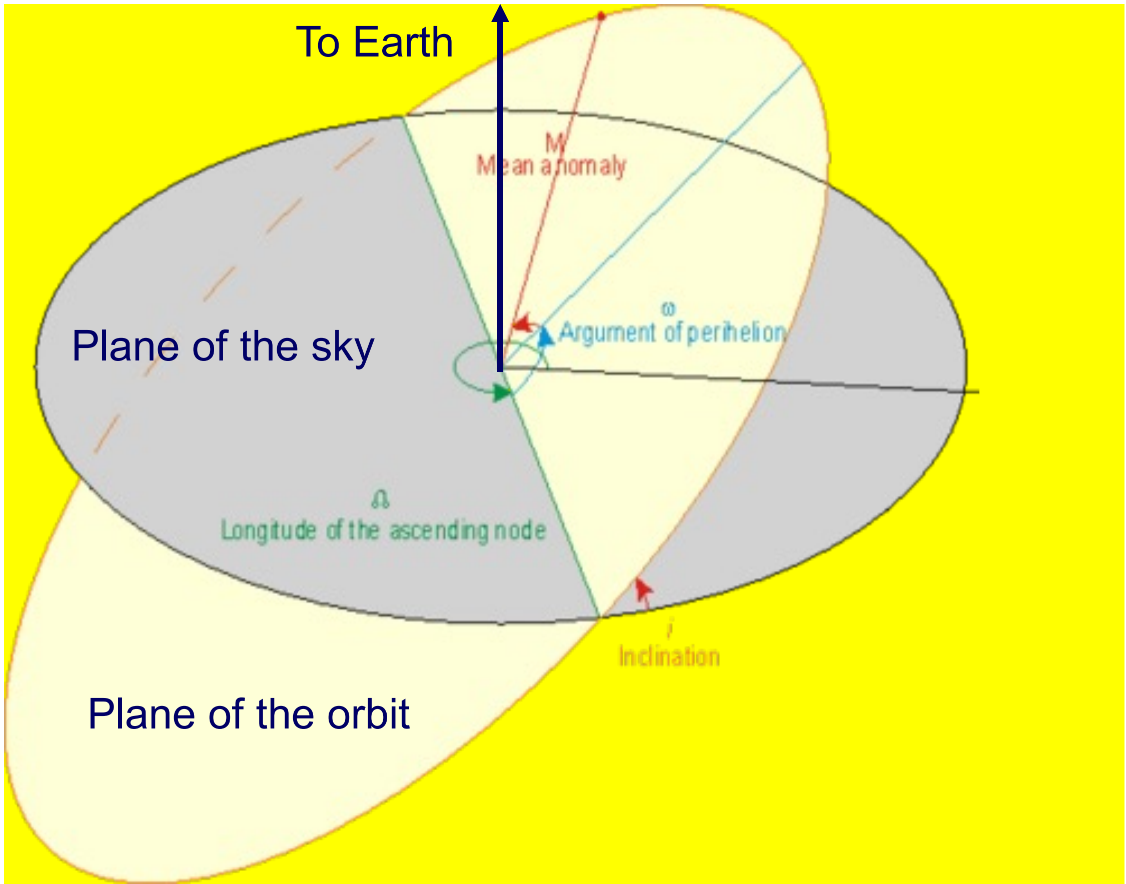

So, whenever you see the residuals like this, you are probably able to find all the keplerian parameters

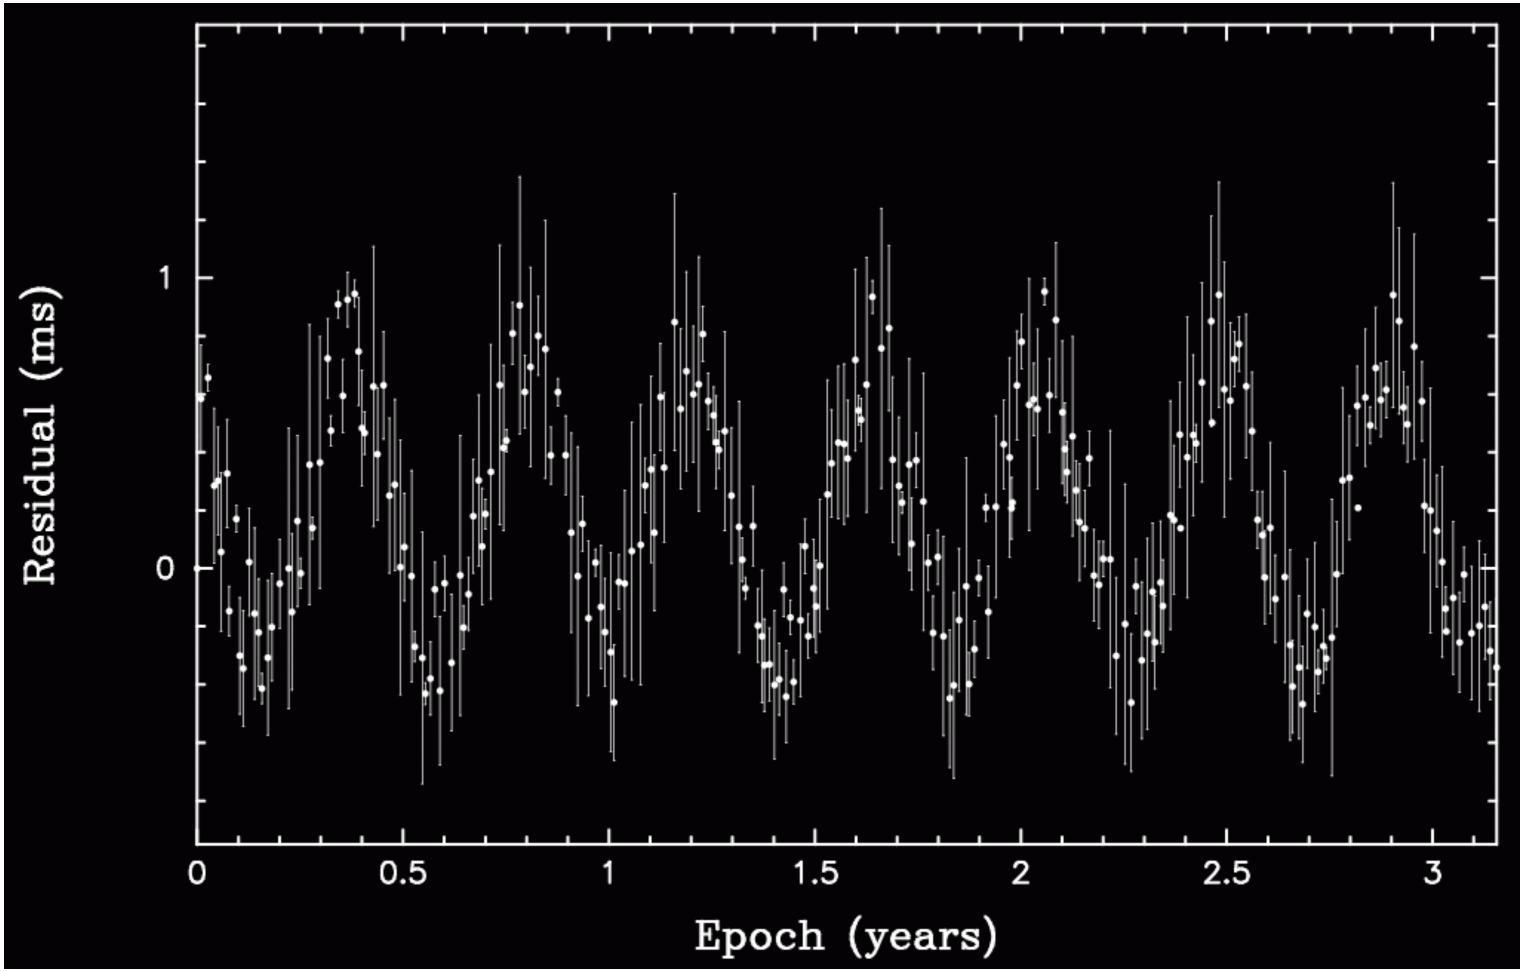

P_b you can measure directly, as well as x. For all the others you have yo analyse the shape of the periodic signal.

All oyu have to do is to translate all the efects, that are going on in the binary into the Roemer delay:

$$\Delta_{\mathrm{RB}} = x(\cos E - e)\sin\omega + x\sin E\sqrt{1-e^2}\cos\omega$$

$$x = \frac{a_p \sin i}{c}$$

E is the eccentric anomaly - one of the variables used to solve the Kepler equation (which describes the movement of a body on an eliptical orbit).

By minimizing the χ^2 function of the residuals you get all the parameters - additionaly to the orbital period you have. The T_0 (time of the periastron passage) is hidden in the *E* value (as it is responsible for the phase of the residuals/orbit).

And of course you cannot tell (usually) the masses of the members of the binary system.

All you have is the mass function:

$$f(m_p, m_c) = \frac{(m_c \sin i)^3}{(m_p + m_c)^2} = \frac{4\pi^2}{G} \frac{(a_p \sin_i)^3}{P_b^2}$$

It allows you to tell only the minimal mass of the companion (and only if you assume that you know the pulsar mass).

So the same mass function you get for a light companion and only slightly inclined orbit, as you get for heavier companion, and the orbit is more inclined.

We see such evaporation of the companion taking place right now - PSR B1957+20. Once per orbital period the pulsar is eclipsed by the matter evaporated from its white dwarf companion.

And because for the white dwarf - the less massive they are - the bigger they get, this will end in a tidal disruption of the companion.

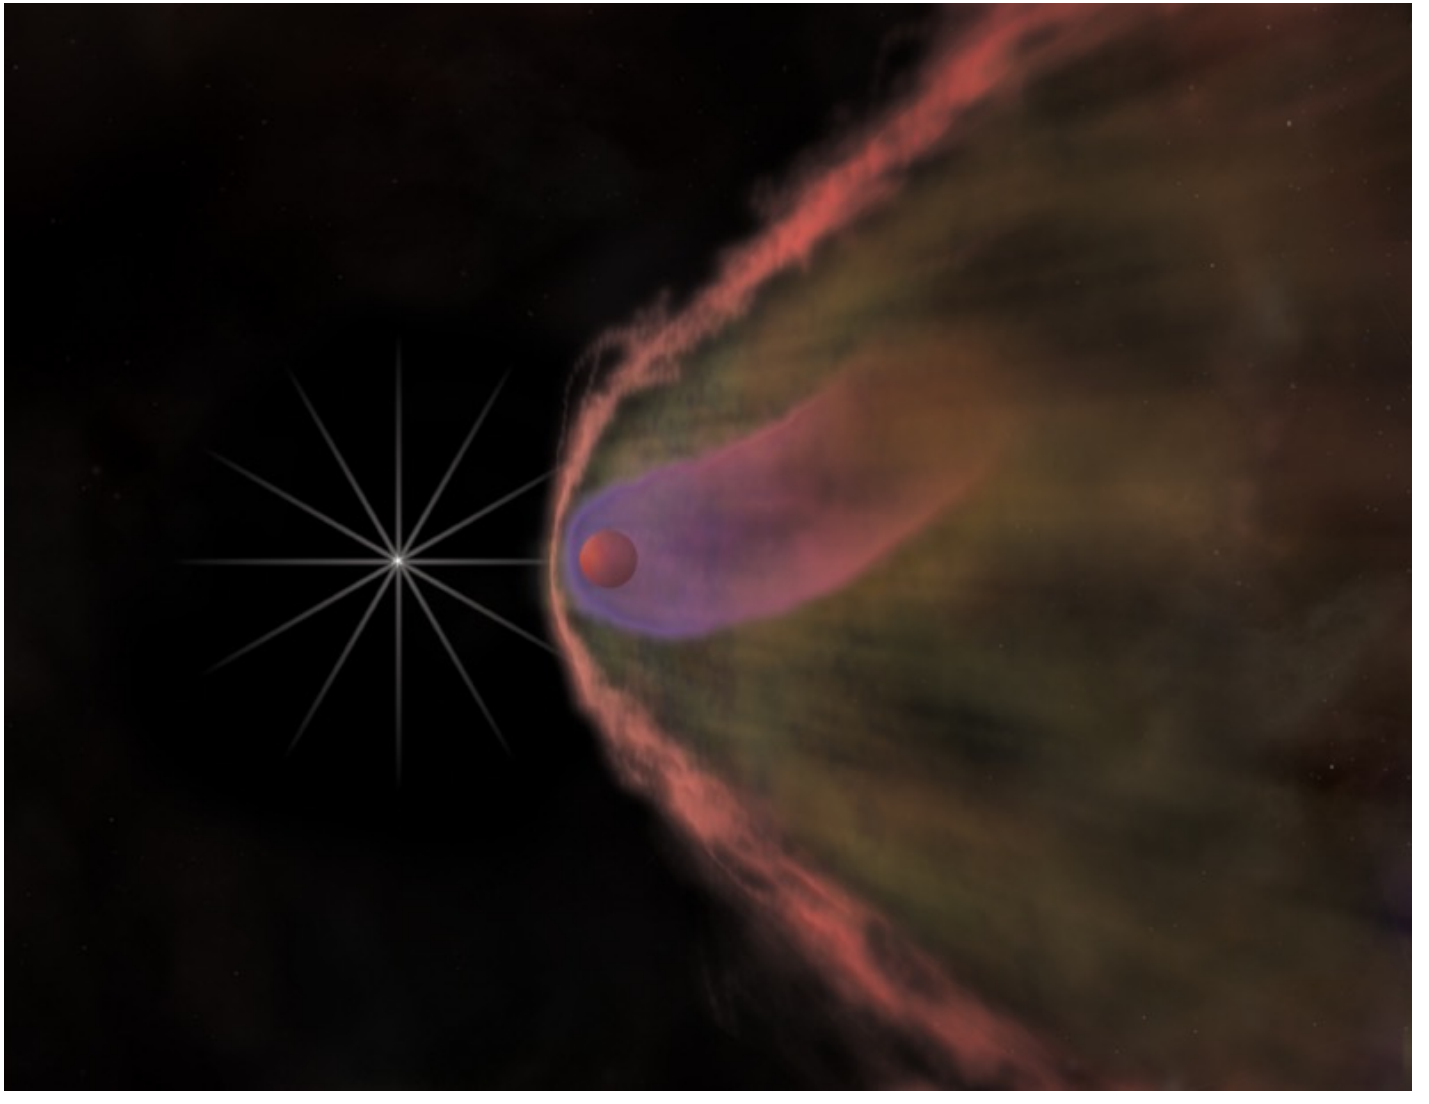

But there are things connected to the binary pulsars, that are even more interesting than planets.

For example - there is a way for us to measure the mass of the companion!

Of course, the Shapiro delay!

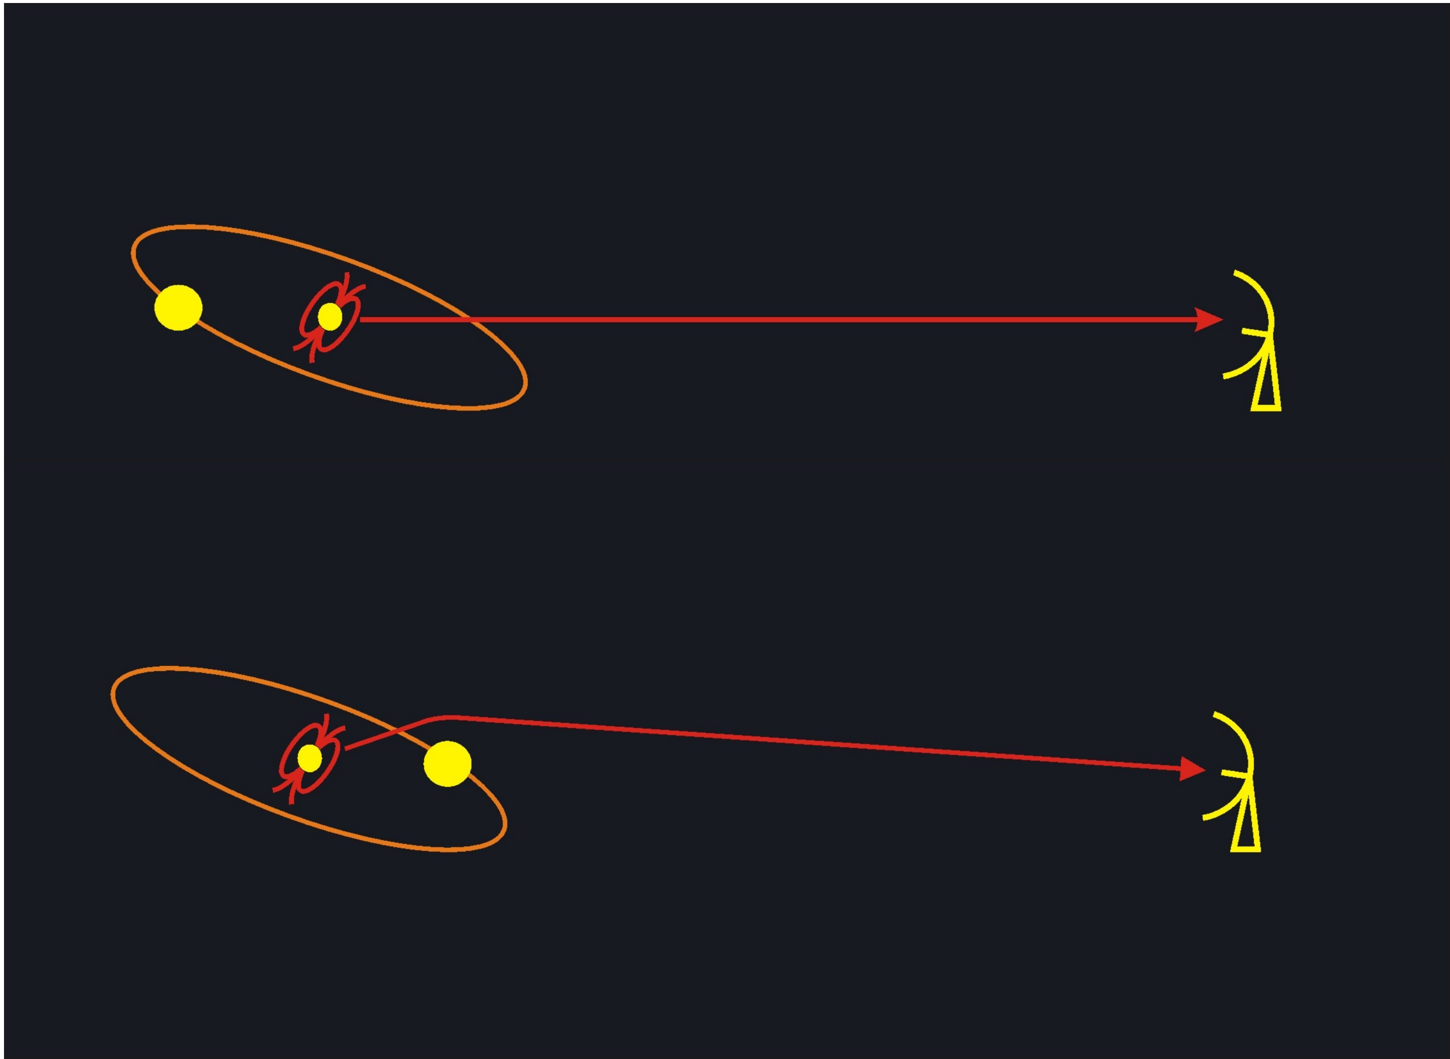

We can measure this effect for seven binary pulsars now. For example the PSR J1640+2224

Pulsar period = 3.16 ms, orbital period = 175 days, eccentricity = 0.0008

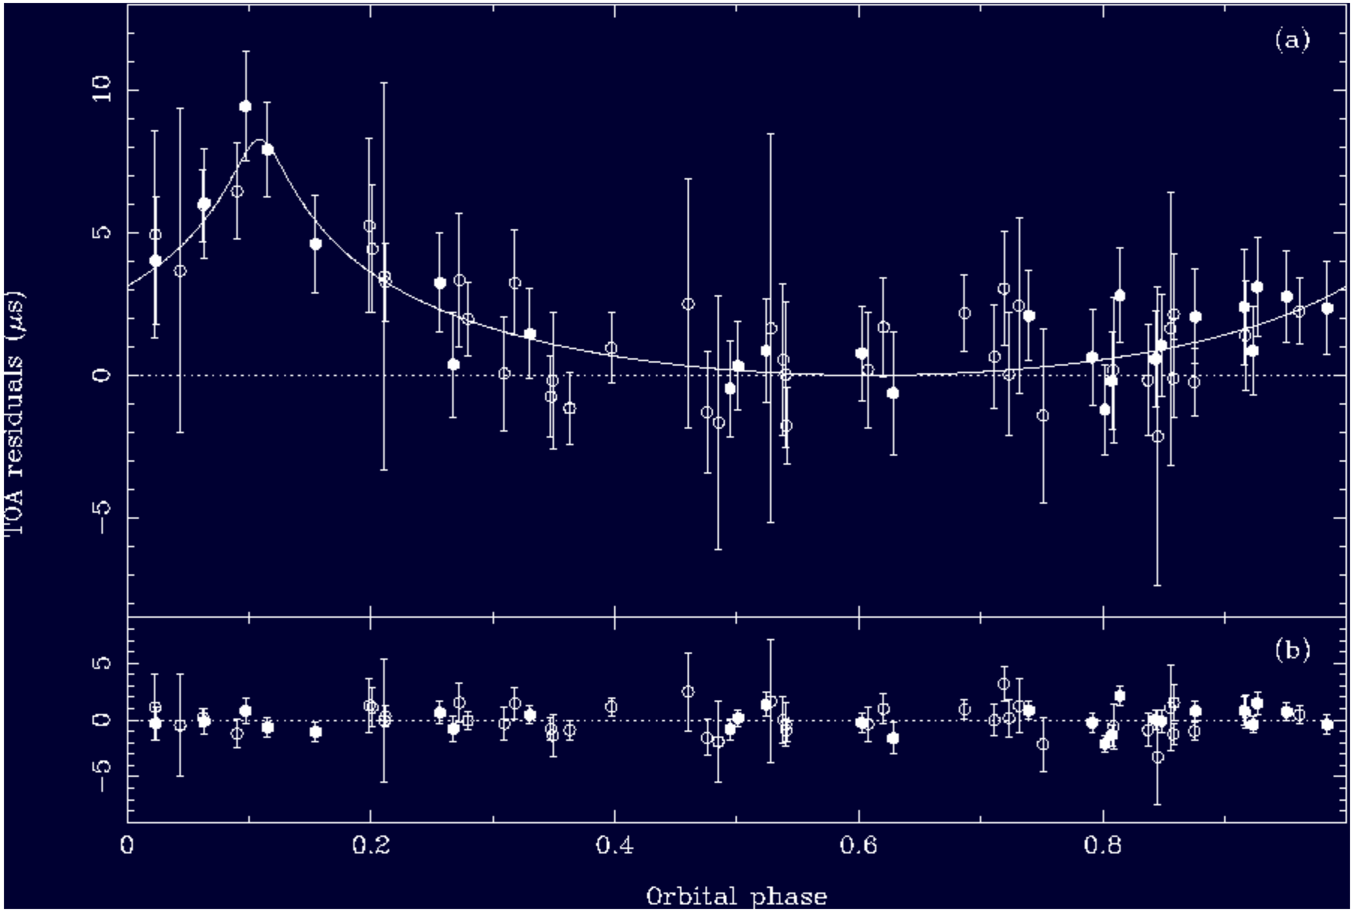

Orbit inclination(degrees) $i = 84^{+4}_{-6}$,

Companion mass(solar masses) $M = 0.15^{+0.08}_{-0.05}$

Translated to the keplerian parameters, the Shapiro delay looks like this:

$$\Delta_{\rm SB} = -2r\ln\left[1 - e\cos E - s\left(\sin\omega(\cos E - e) + \sqrt{1-e^2}\cos\omega\sin E\right)\right]$$

Where the parameter *r* (range) depends solely on the companion mass, and s (shape) is just proportional to sin *i*.

Taken from "Handbook of Pulsar Astronomy" by Lorimer & Kramer

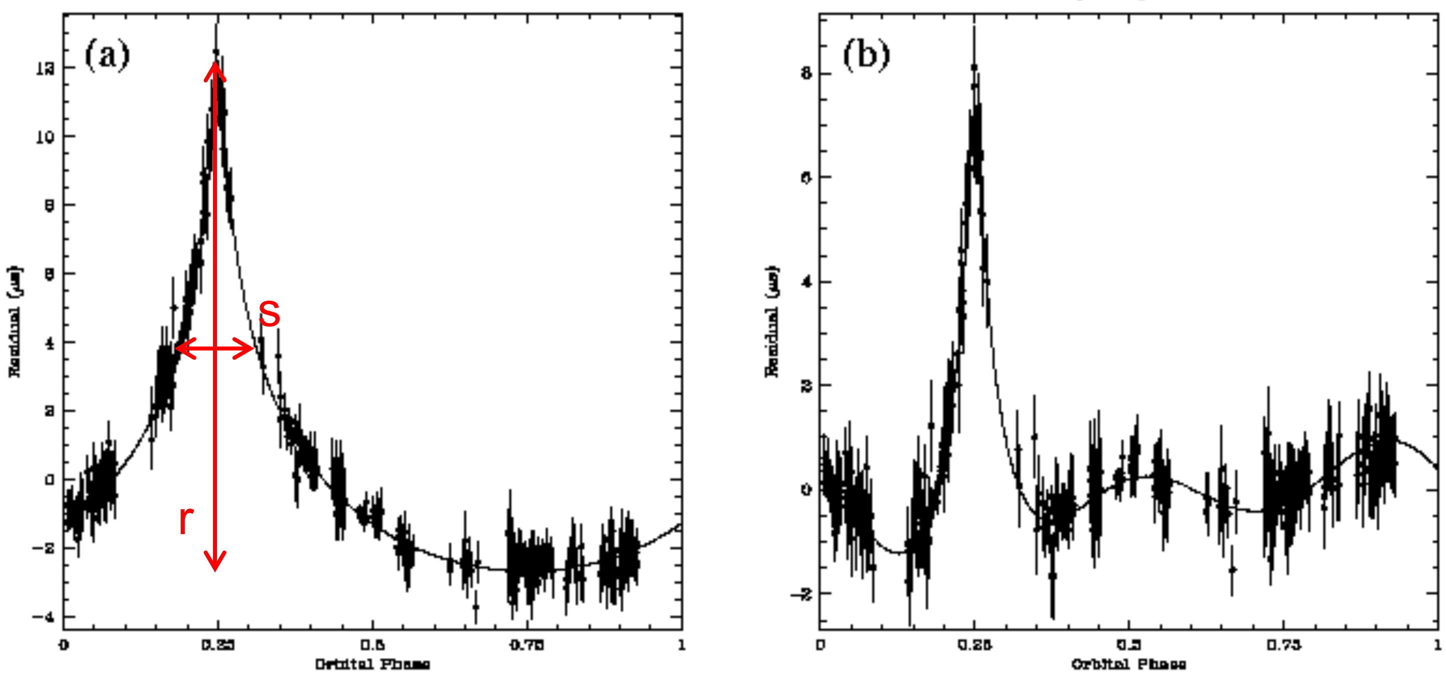

Of course, if you can measure the keplerian parameters, you can also tell if they change in time.

One of such effects is the periastron advance - a phenomenon, which is visible in our solar system - the peihelion advance of Mercury.

In the case of Mercury it is only 43 arcsecond per century. For tight pulsar binaries this effect can be measured easily over a period of a few years.

In terms of keplerian parameters it is seen as the change of the periastron longitude - ω.

This is one of the effects forseen by the General Relativity.

.

The periastron advance is given by:

$$\dot{\omega} \propto \left(\frac{GM}{ac^2}\right)\frac{1}{1-e^2}$$

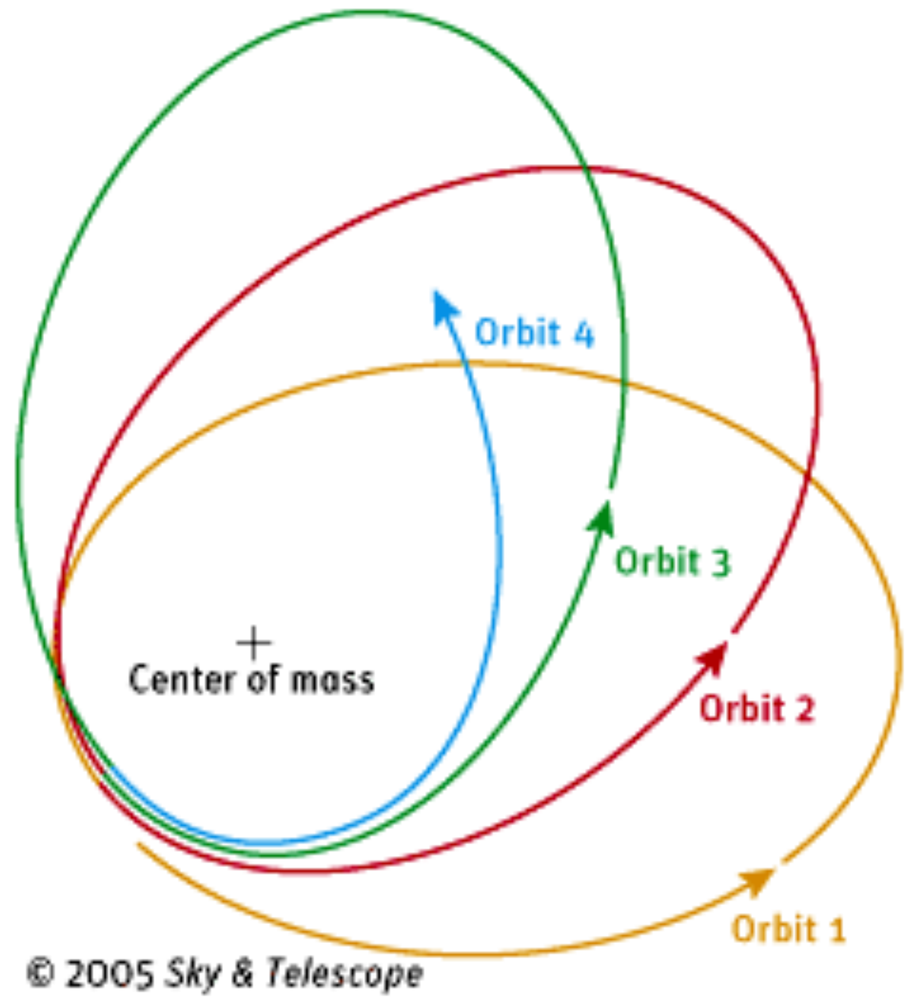

But it is not enough to say, that Einstein was right, as some other post-newtonian theories of the gravity can explain that.

But the shortening of the orbital period - it is another thing.

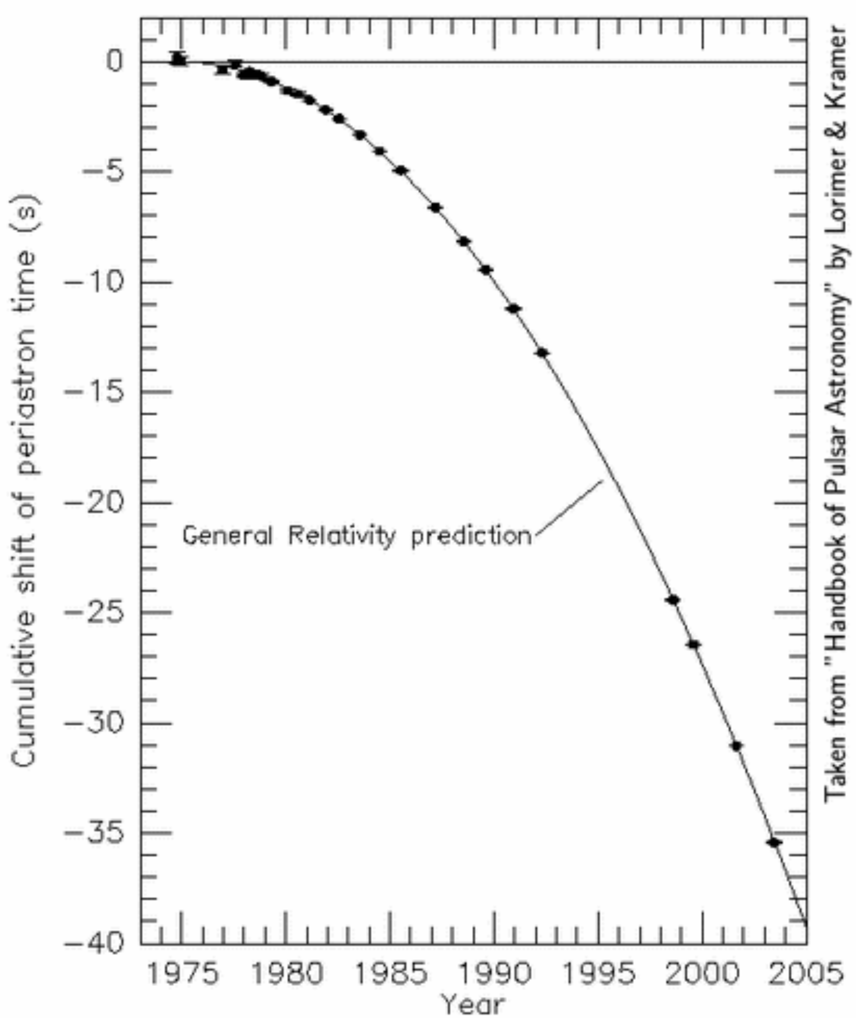

Measutement of the P_b in the pulsar B1913+16 - a double neutron star system - showed it is going exactly as the Einstein theory predicted.

For the confirmation of Einstein's General Relativity R.Hulse and J.Taylor were awarded The Nobel Prize in Physics in 1993.

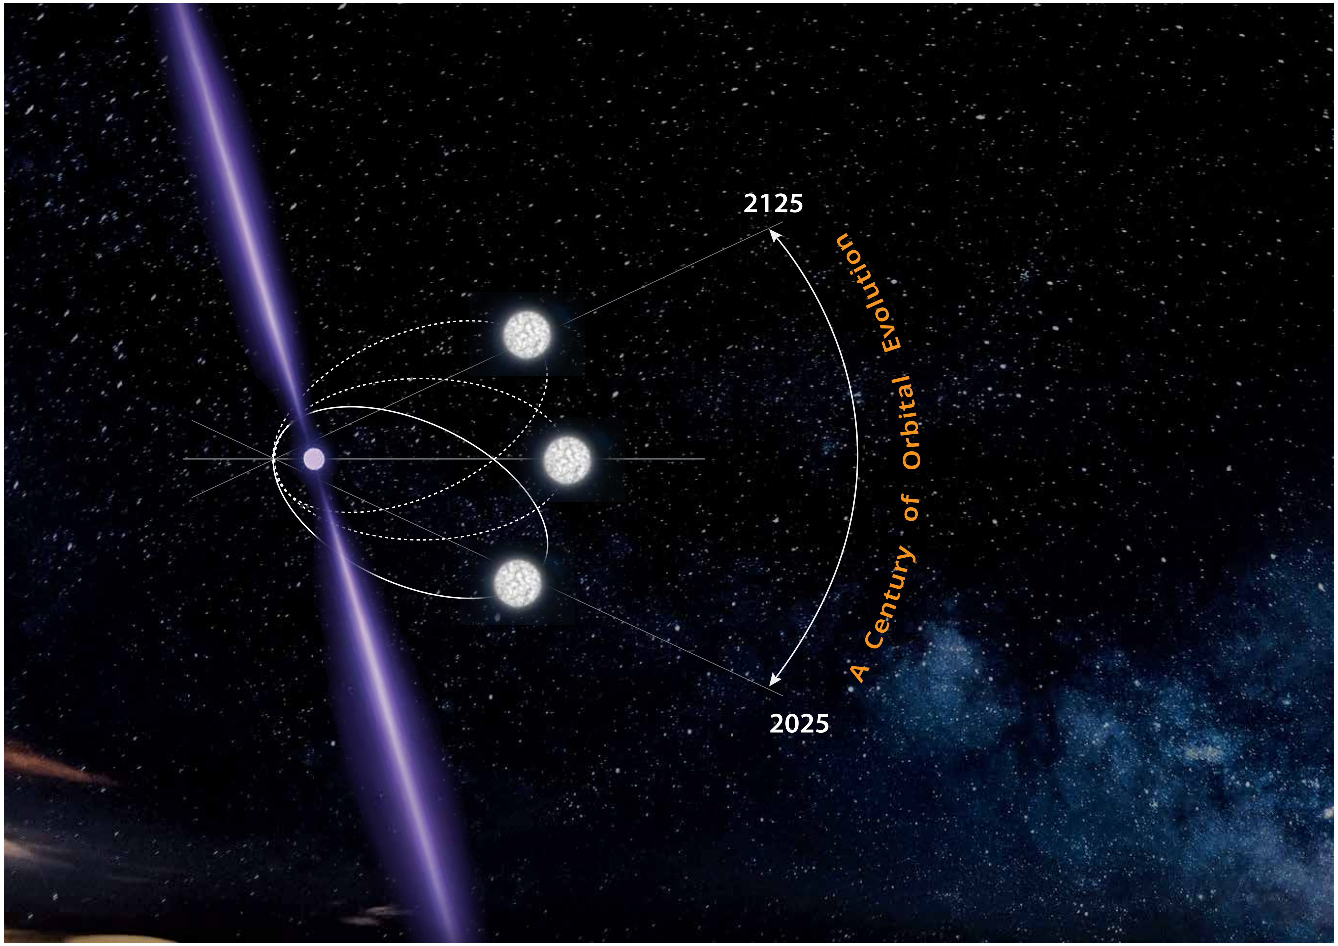

J1617–2258A discovered in M80 Globular cluster by Jyotirmoy (Das et al. 2025)

This precision is about half a degree per year. To put that in perspective, this pulsar's orbit shifts in a single day by roughly as much as Mercury's perihelion shifts in an entire decade." That precise measurement lets astronomers 'weigh' the system. Together, the pulsar and its companion have a mass of about 1.67 times that of the Sun.

If the orbital period is shortening this means, that the system is actualy shrinking (so we can measure *x*). According to Einstein's theory this is due to the emission of gravitational waves:

projected light travel time across the orbit

$$x = \frac{a_p \sin i}{c}$$

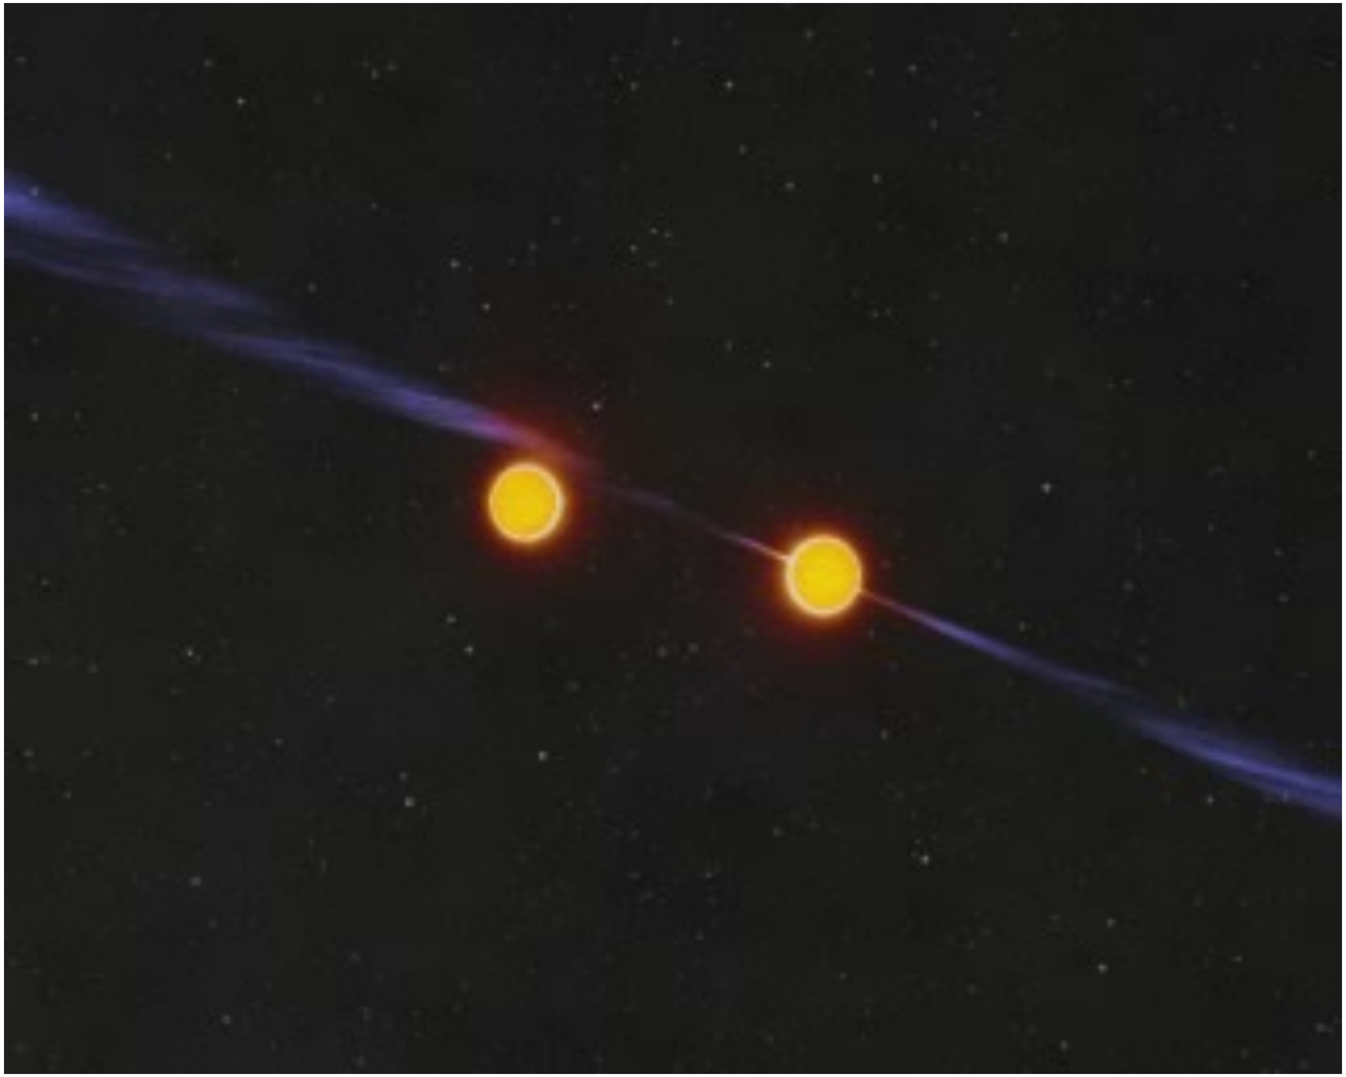

Movie credit:
John Rowe
Animations

But let's go back to our pulsars.

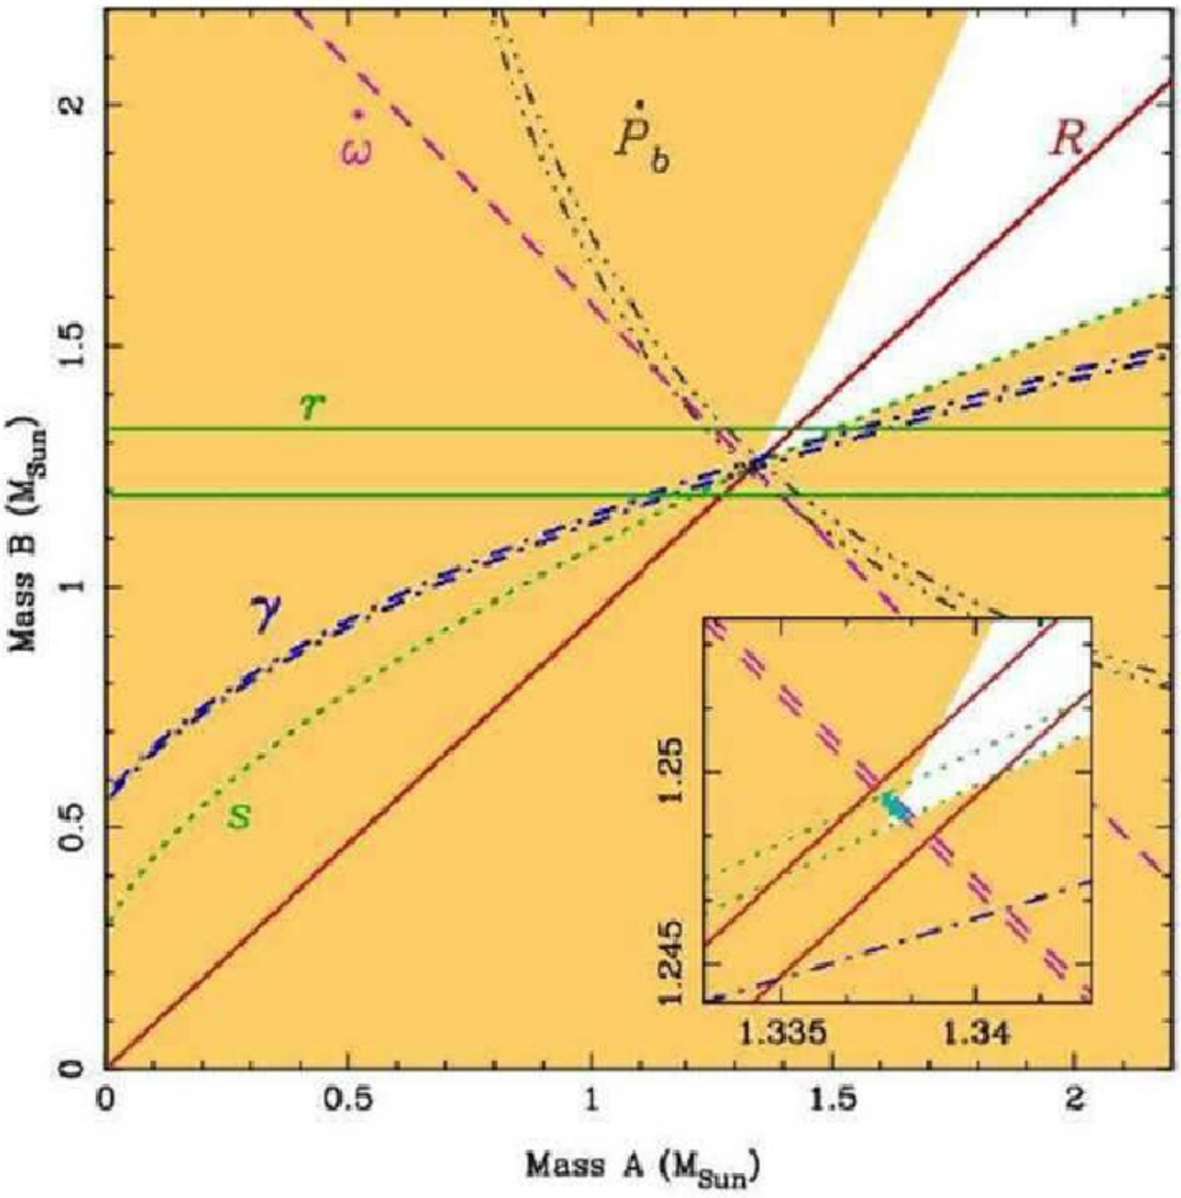

Thanks to the observations of a double neutron star binaries we can put limits on their masses. Using the Shapiro delay, the measyrements of ω and P_b, and the effect of grawitational redshift (Einstein delay, represented by the γ parameter) we can tell them quite precissely...

Parameter	Depends on
$\dot{\omega}$	total mass
γ	both masses
$\dot{P}_b$	masses + GW emission
r	companion mass
s	inclination

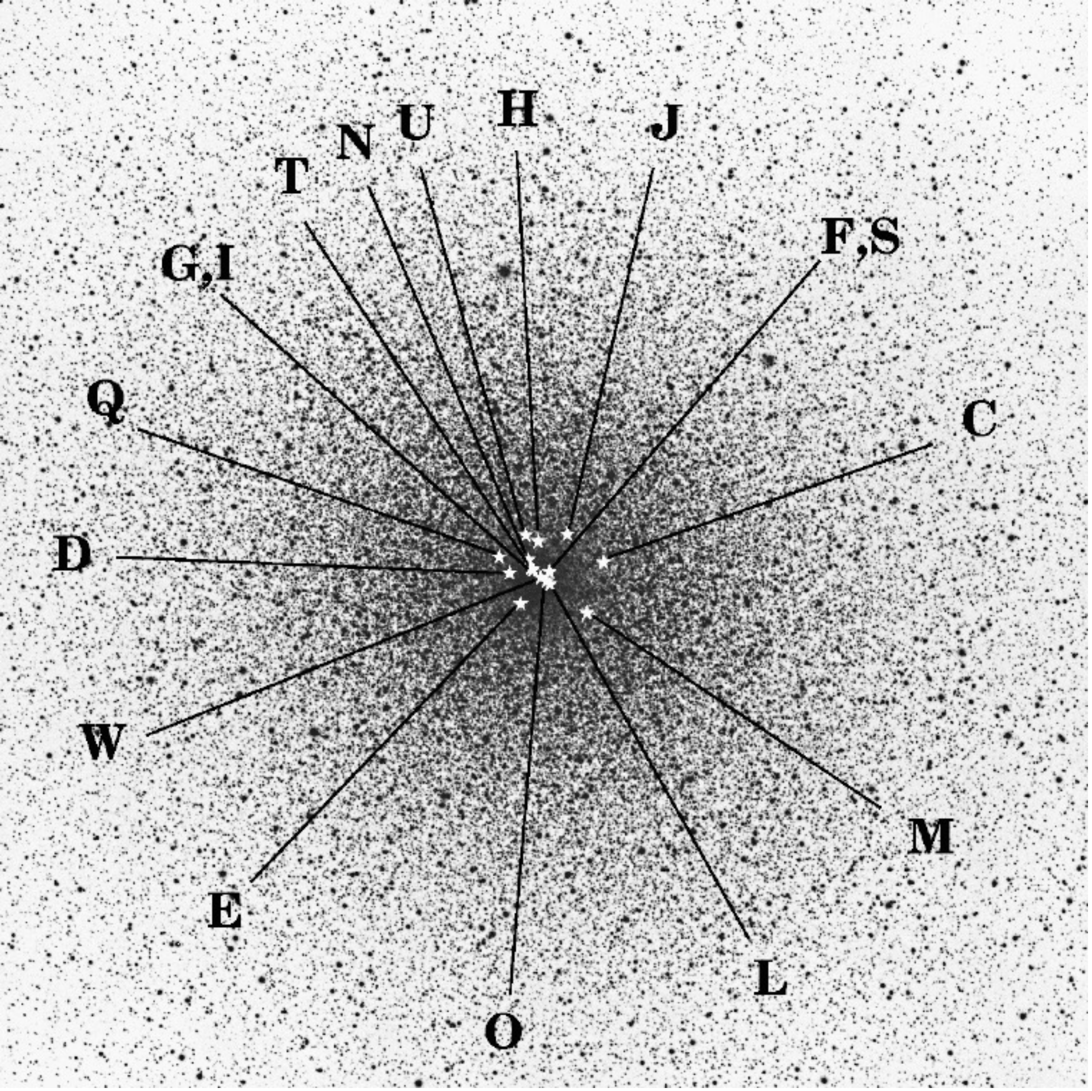

http://www2.naic.edu/~pfreire/47Tuc/

In the cluster 47 Tucanae we found 23 pulsars so far. And still counting...

#	NAME	P0 (s)	P1
1	B0021-72C	0.005757	-4.98e-20
2	B0021-72D	0.005358	-3.43e-21
3	B0021-72E	0.003536	9.85e-20
4	B0021-72F	0.002624	6.45e-20
5	B0021-72G	0.004040	-4.22e-20
6	B0021-72H	0.003210	-1.83e-21
7	B0021-72I	0.003485	-4.59e-20
8	B0021-72J	0.002101	-9.79e-21
9	B0021-72L	0.004346	-1.22e-19
10	B0021-72M	0.003677	-3.84e-20
11	B0021-72N	0.003054	-2.19e-20
12	J0024-7204O	0.002643	3.04e-20
13	J0024-7204P	0.003643	*
14	J0024-7204Q	0.004033	3.40e-20
15	J0024-7204R	0.003480	*
16	J0024-7204S	0.002830	-1.21e-19
17	J0024-7204T	0.007588	2.94e-19
18	J0024-7204U	0.004343	9.52e-20
19	J0024-7204V	0.004810	*
20	J0024-7204W	0.002352	*

ATNF pulsar cataogue

When doing timing, or any other type of pulsar observations...

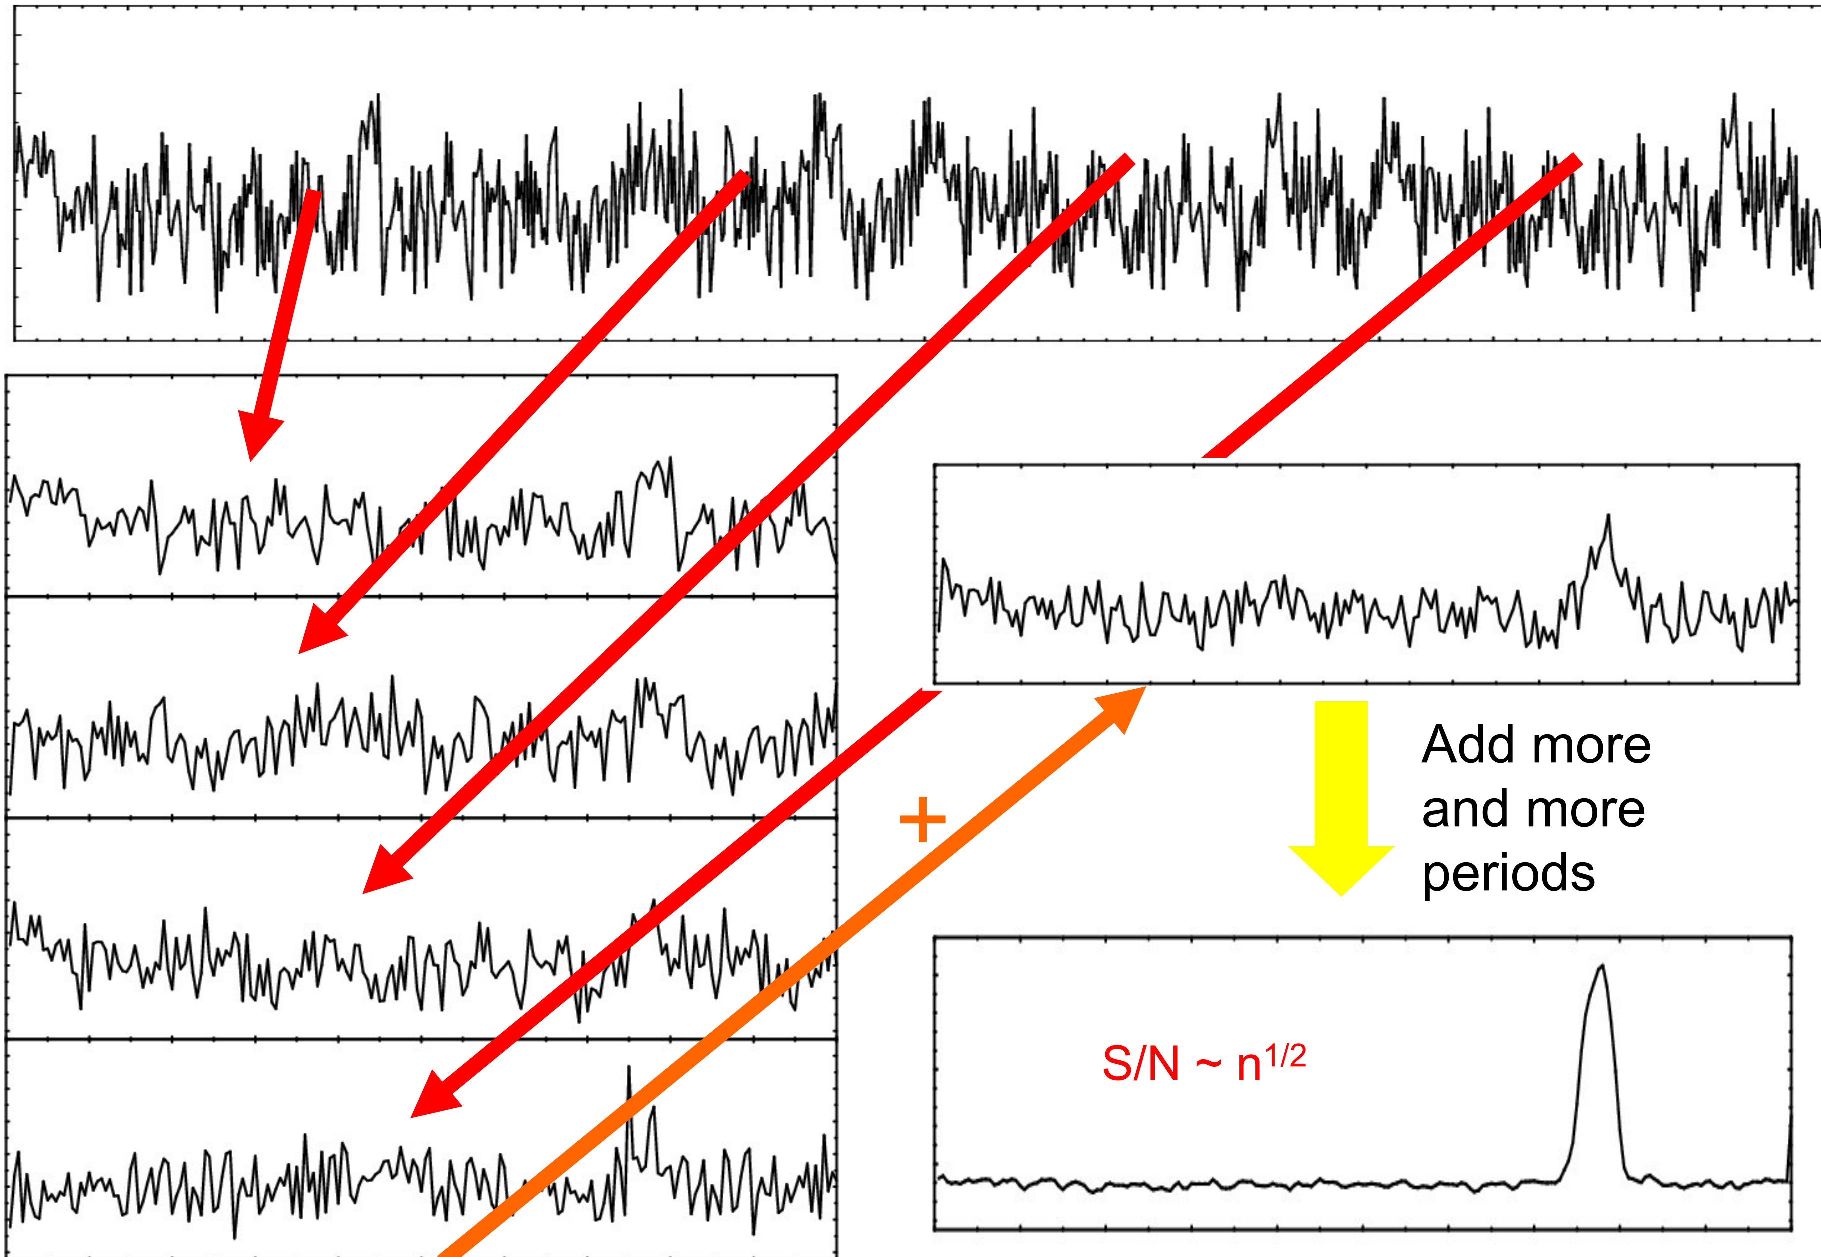

But for searches, we just have the time series and no knowledge about the pulsar period, right?

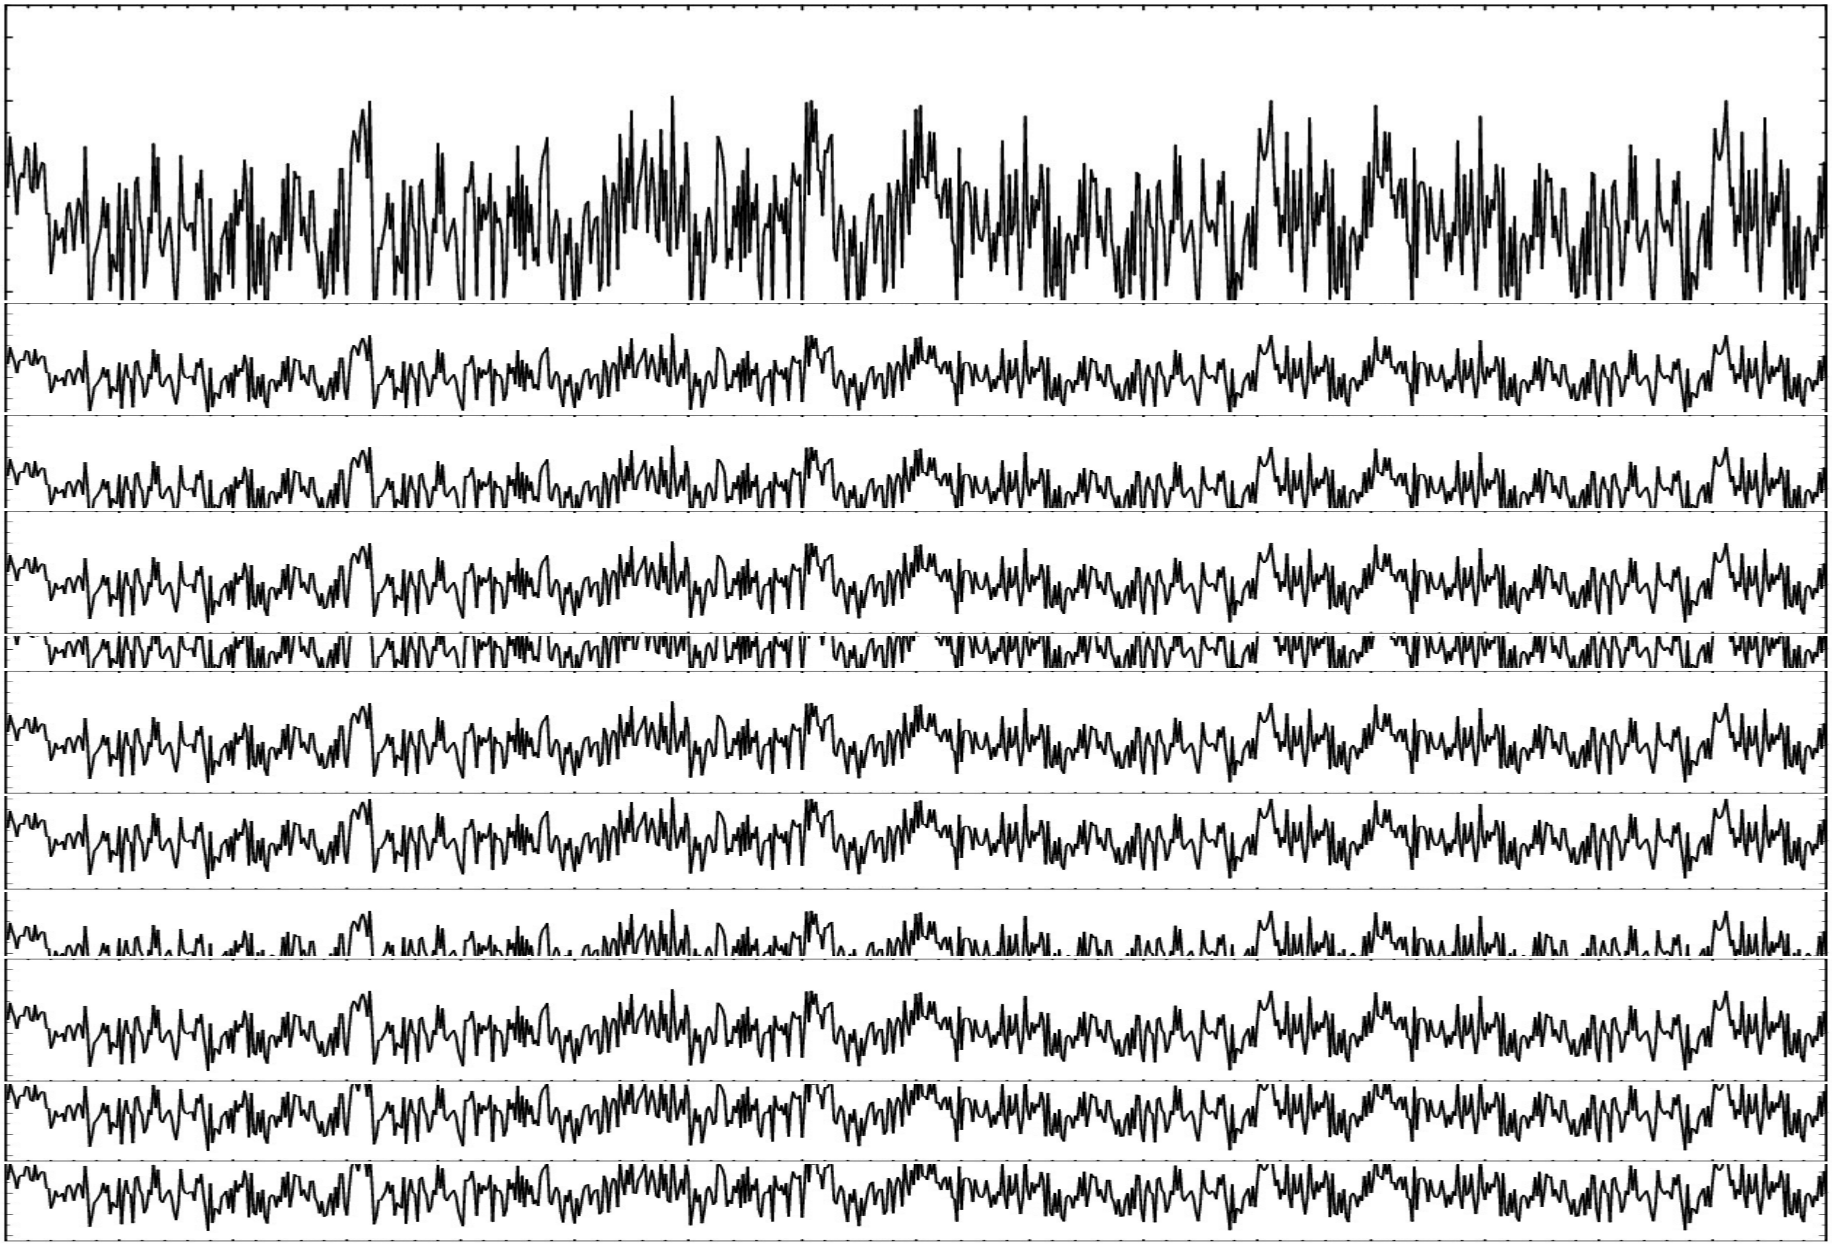

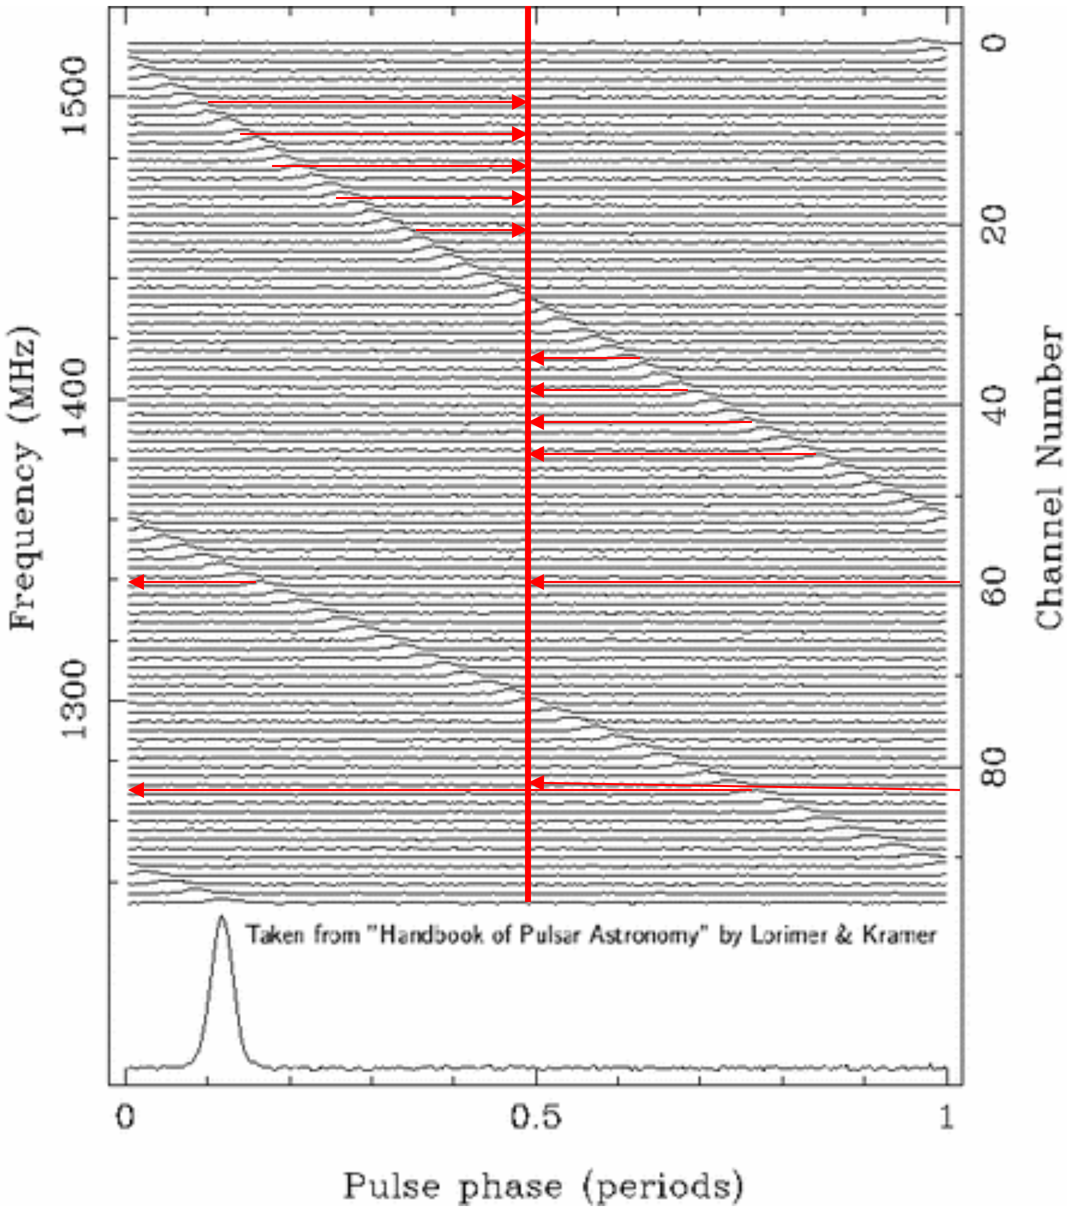

You have to dedisperse your time series.

And since you don't know the disspersion measure of a pulsar you will find (usually), you have to try a latge number of possible "trial" DM's.

Basically in the simpliest case to find a pulsar, you have to perform a 2D search in the P - DM plane.

Signal to noise ratio of the pulsar depends on the difference between the real DM and you trial DM, but it also depends on the pulsar period.

To find millisecond pulsars the grid of the trial DM values has to be very dense.

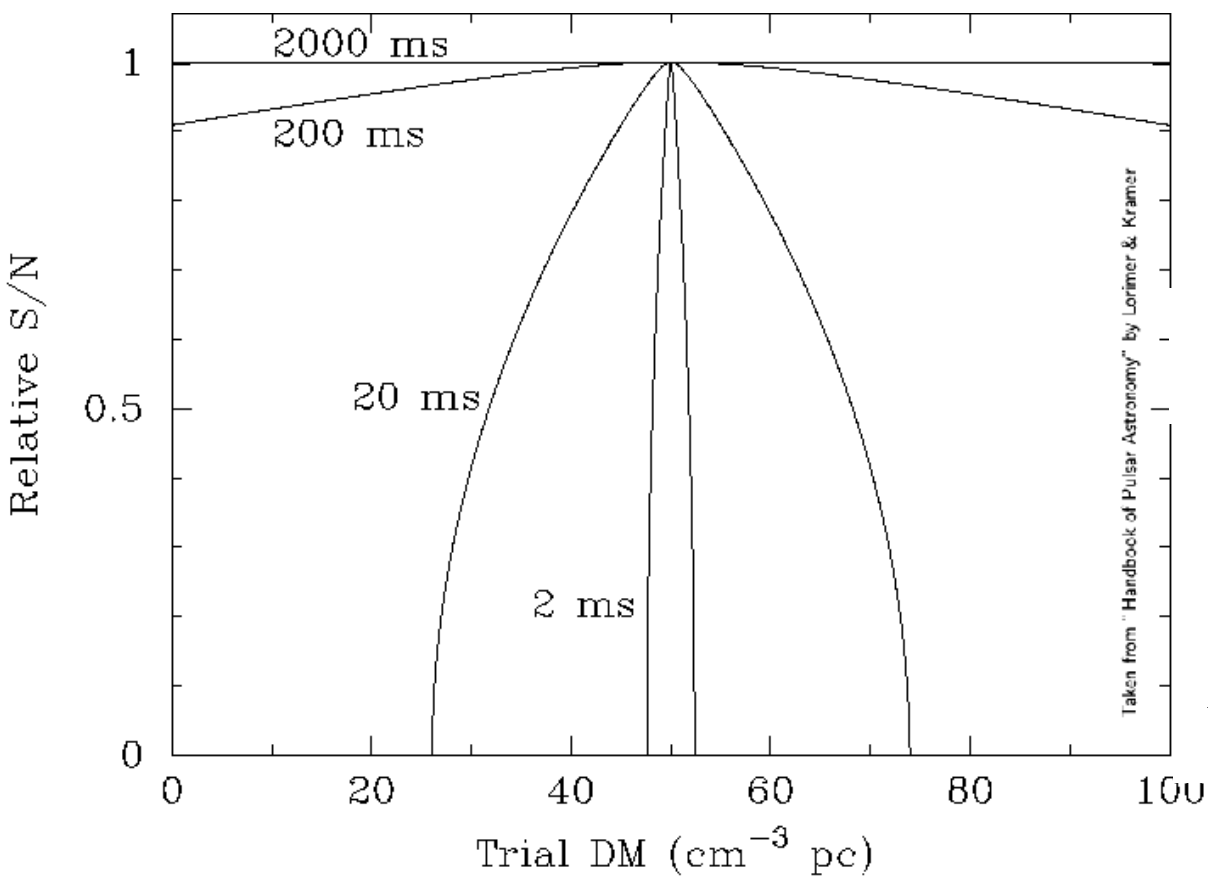

It is exceptionally hard to find millisecond puslars with high DM values.

$$t_{\rm res} = k_{\rm DM}\,|\Delta{\rm DM}|\,\left(\nu_{\rm low}^{-2} - \nu_{\rm high}^{-2}\right)$$

$$W_{\rm obs} \simeq \sqrt{W^2 + t_{\rm res}^2}$$

$$\frac{({\rm S/N})_{\rm trial}}{({\rm S/N})_{\rm max}} \simeq \sqrt{\frac{W}{W_{\rm obs}}} = \left(1 + \frac{t_{\rm res}^2}{W^2}\right)^{-1/4}$$

This is because the DM discrepancy broadens (smears) the pulse over a larger range of pulse phases, than its intrinsic values:

$$W_{\rm eff} = \sqrt{W_{\rm int}^2 + (k_{DM} \times |DM| \times \Delta f / f^3)^2}$$

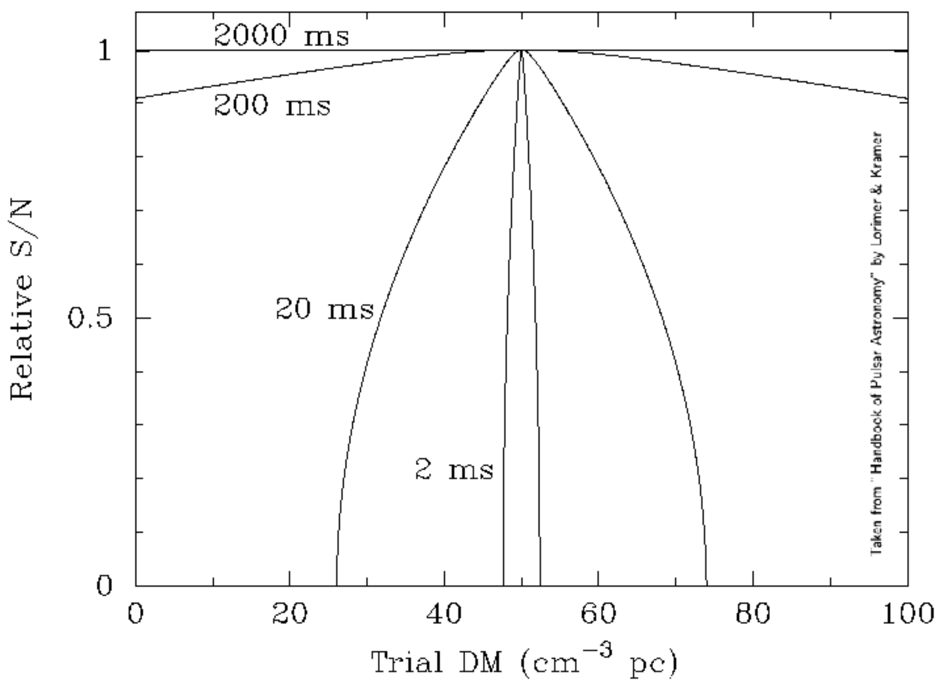

$$S/N \propto \sqrt{\frac{P - W_{\rm eff}}{W_{\rm eff}}}$$

The lower the frequency is - the more prominent is the effect, and one needs more trial DM's.

Amounts of data to analyse are tremendous. Exmaple (from the PSPM):
at sampling rate of 100 ms 28 min integraton means 2^{24} samples.
That is: 2^{24} samples in each of 128 spectral channels. That makes 2^{31} values altogether (1 GB of data in the case of PSPM).
And you have usually a few hundred "trial DM" values to check...

The dedispersion of the time series

$$\mathcal{T}_j = \sum_{l=1}^{n_{ch}} \mathcal{R}_{j+k(l),l} \ ,$$

$$k(l) = \left(\frac{t_{\rm samp}}{4.15 \times 10^6 {\rm ms}}\right)^{-1} \left(\frac{DM}{\rm cm\ pc^{-3}}\right) \left[\left(\frac{f_l}{\rm MHz}\right)^{-2} - \left(\frac{f_0}{\rm MHz}\right)^{-2}\right]$$

$$f_l = f_0 - (l-1)\Delta f$$

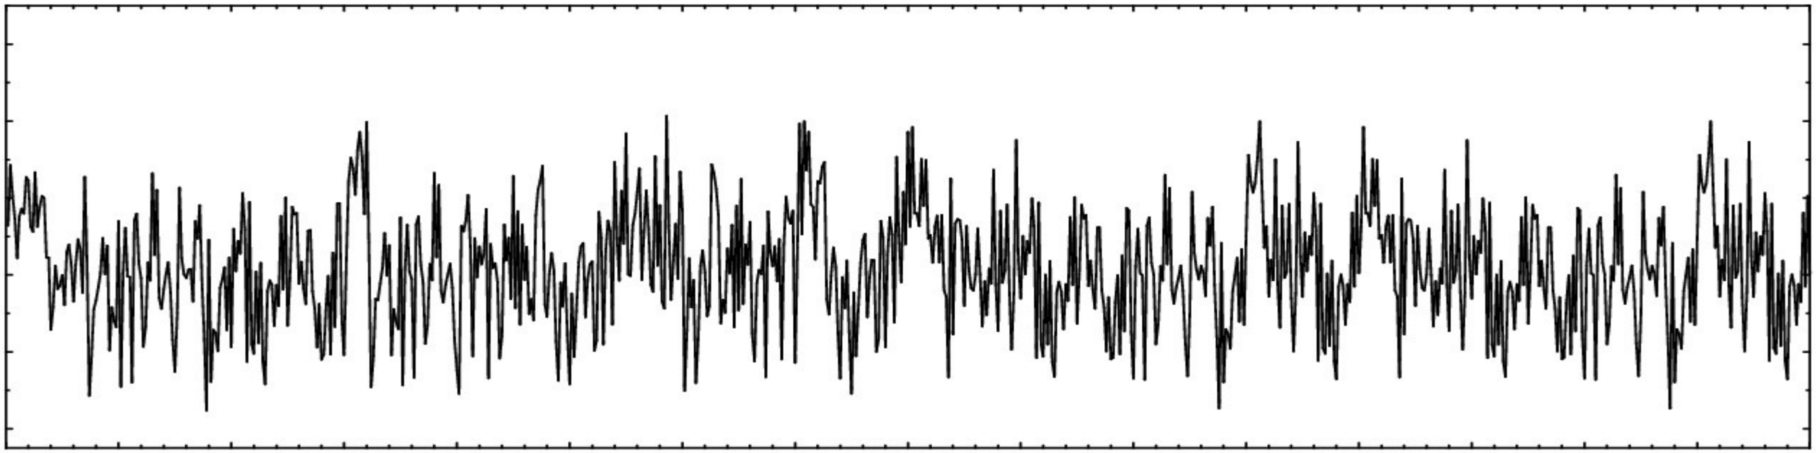

Assume, that you have the timeseries at one trial DM.

Most of the methods of finding periodicity in a dataset are Fourier-based methods, and probably the mostcommon is the FFT (Fast Fourier Transform).

$$\mathcal{F}_k = \sum_{j=0}^{N-1} \mathcal{T}_j \exp(2\pi i j k/N),$$

As a result of FFT you'll have a set of Fourier transform values at N/2 frequencies (where N is the number of samples in your timeseries);

$$\nu_k = \frac{k}{N t_{\mathrm{samp}}} = \frac{k}{T} \qquad 1 \leqslant k \leqslant N/2$$

The lag between the frequencies depends purely on the lenght of your timeseries.

$$\Delta\nu_k = \frac{1}{T}$$

The range of the frequencies is limited by the total timespan, and the Nyquvist frequency.

$$\frac{1}{T} < \nu_k < \frac{1}{2t_{\mathrm{samp}}} = \nu_{\mathrm{Nyq}}$$

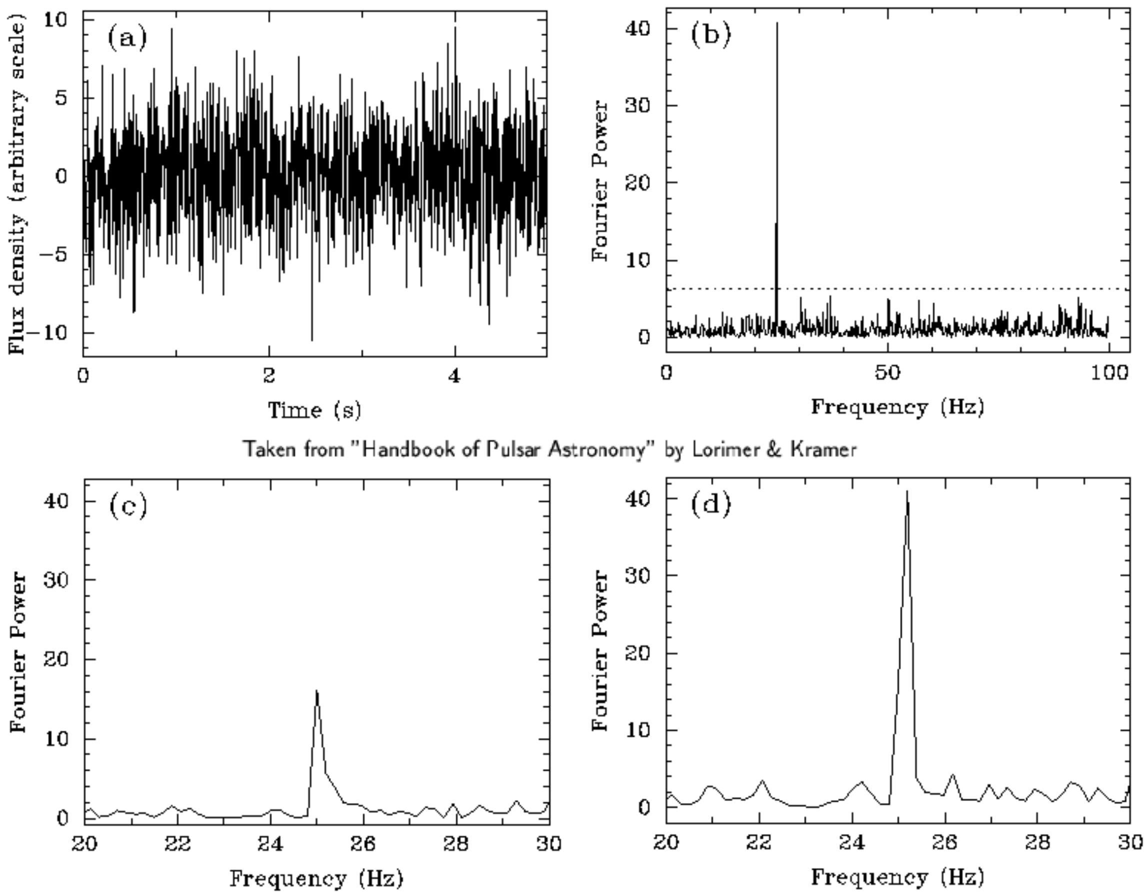

Sample timeseries [a], Fourier power spectrum [b], a close-up of one of the candidates [c], and the corrected fourier power [d].

If the real frequency of the signal falls in between your fourier bins you may lose up to 60% of its power!

$$\mathcal{F}_{k+\frac{1}{2}} \simeq \frac{\pi}{4}(\mathcal{F}_k - \mathcal{F}_{k+1})$$

$$\mathcal{P}_{k+\frac{1}{2}} = |\mathcal{F}_{k+\frac{1}{2}}|^2$$

Some "recovery" methods are required, but usually the simple ones suffice - as for example the Fourier-domain interpolation (loss of power is limited to 7% at most.

Correcting for the low frequency (red) noise, due to receiver system instabilities, the weather and the RFI.

It's important especially for finding long-perod pulsars.

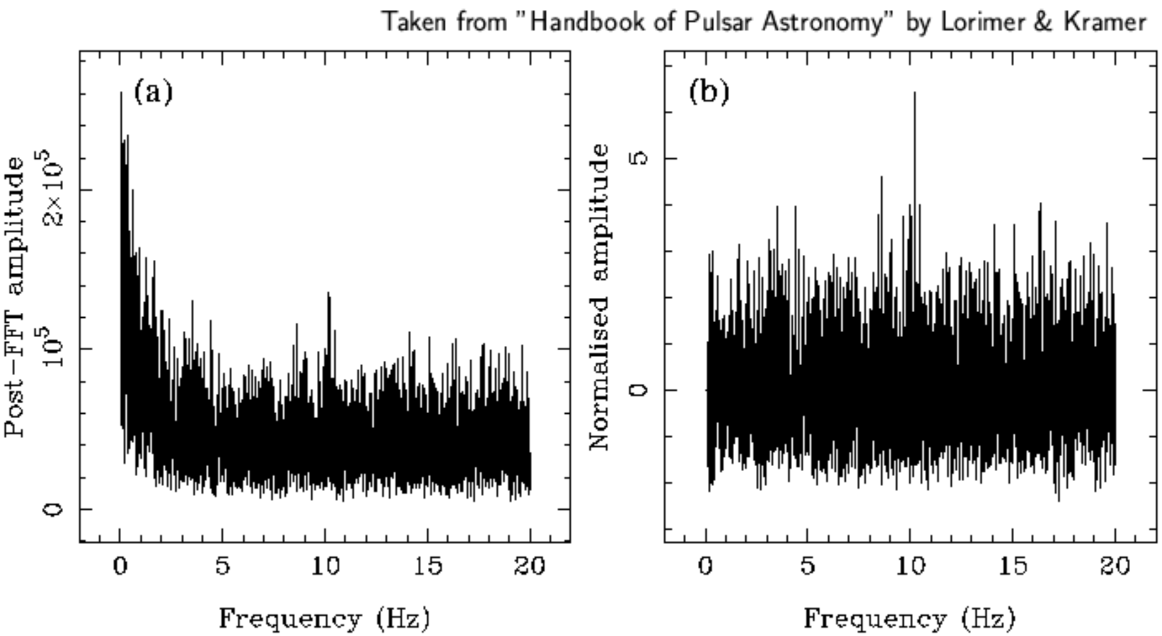

One has to apply some "whitenning" algorithms.

But usually that is not the case ;(

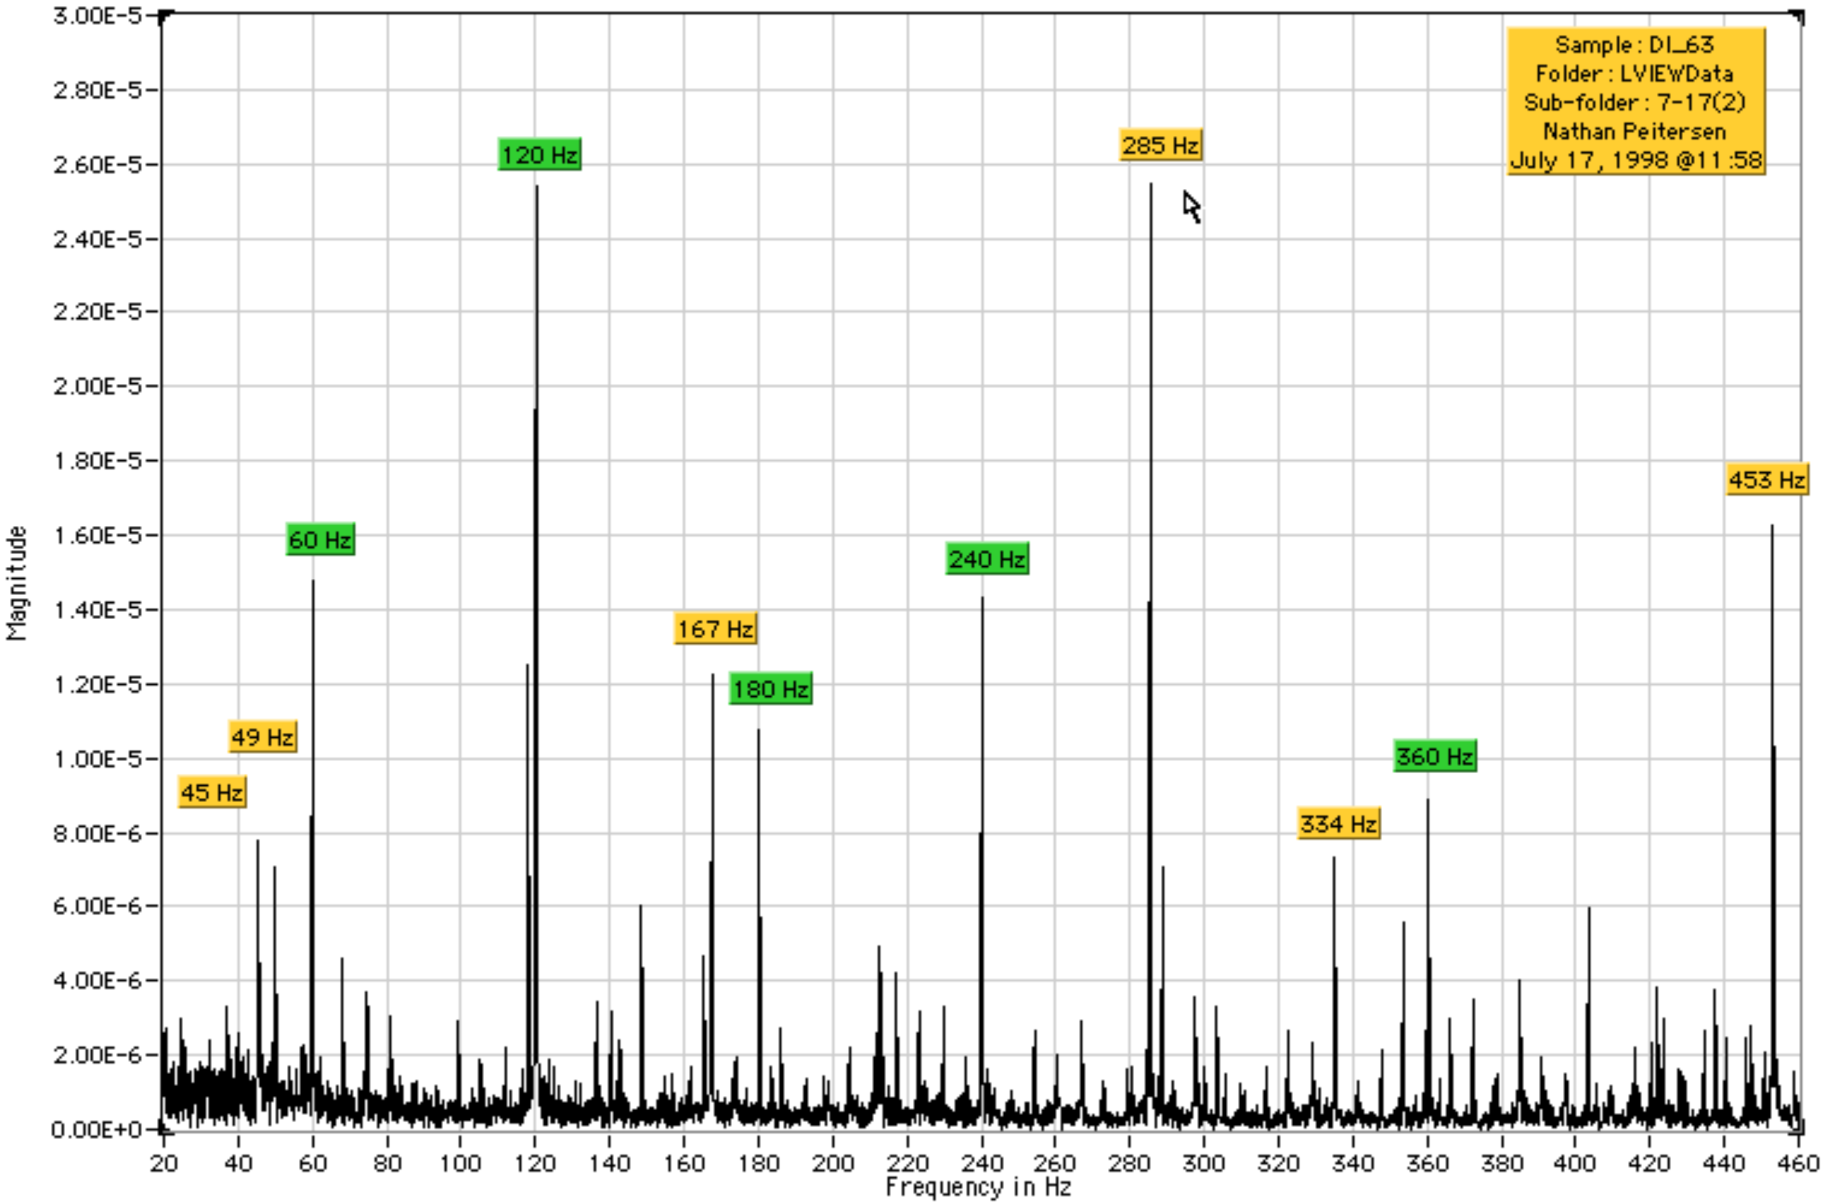

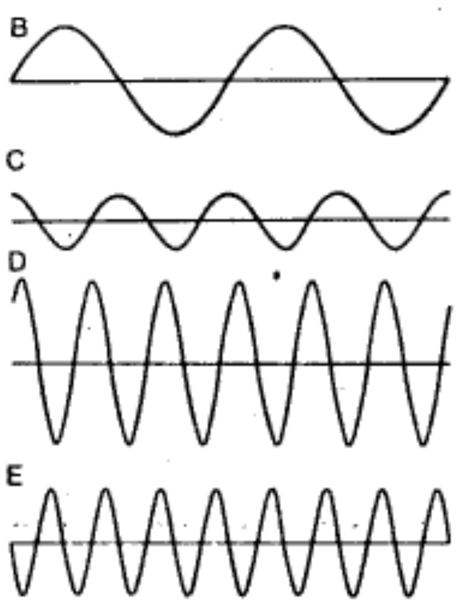

Harmonics:

• their strength depends on the shape of the profile

• for a sinewave they "die" quickly

• the less the duty cycle is, the more power in the harmonics you have

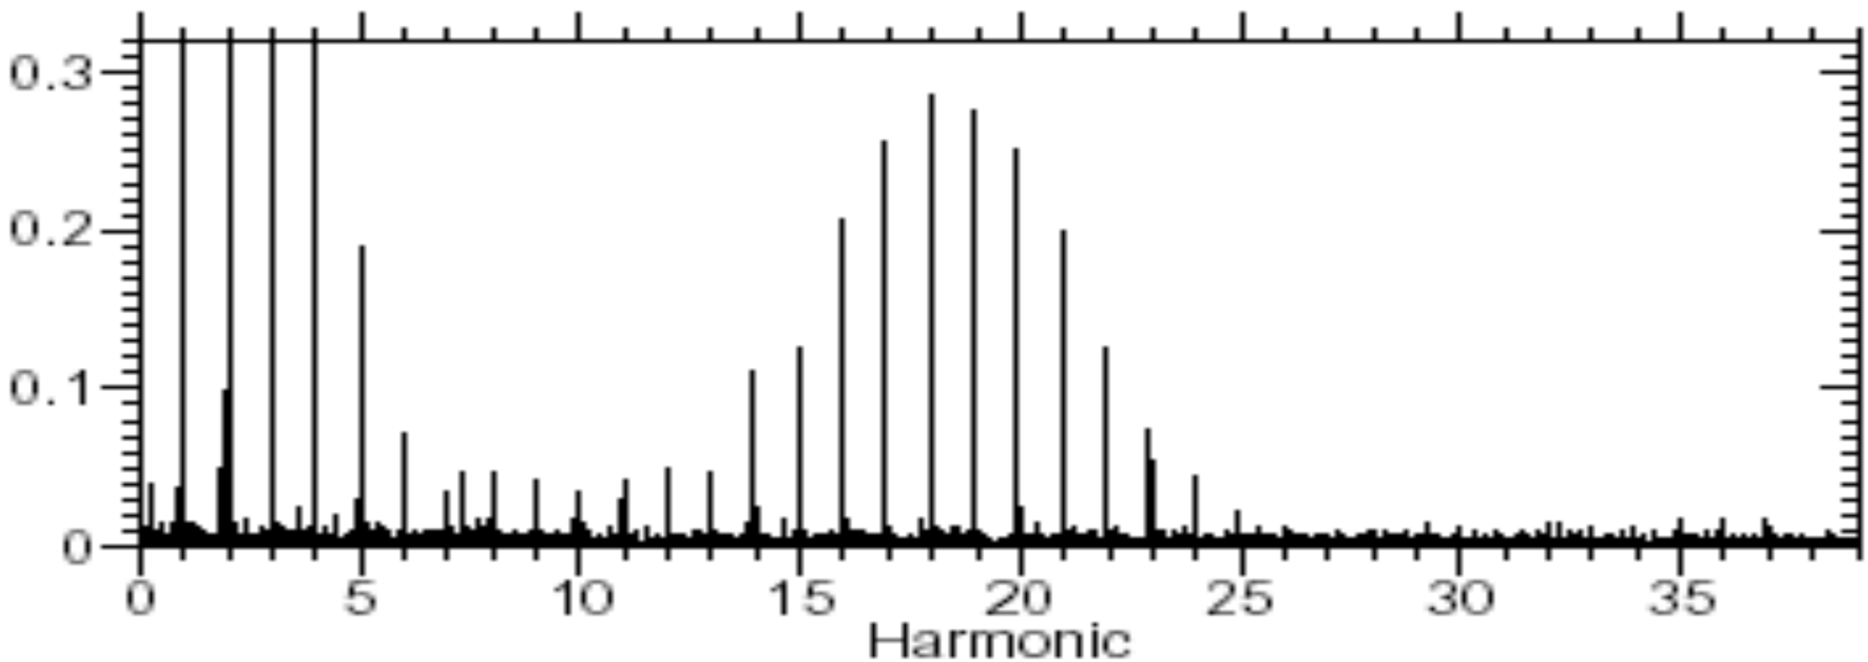

You can use that to your advantage!

Harmonic summing: noise increases as 2 1/2, but power can increase by a factor of 2 - this means 2 1/2 gain!

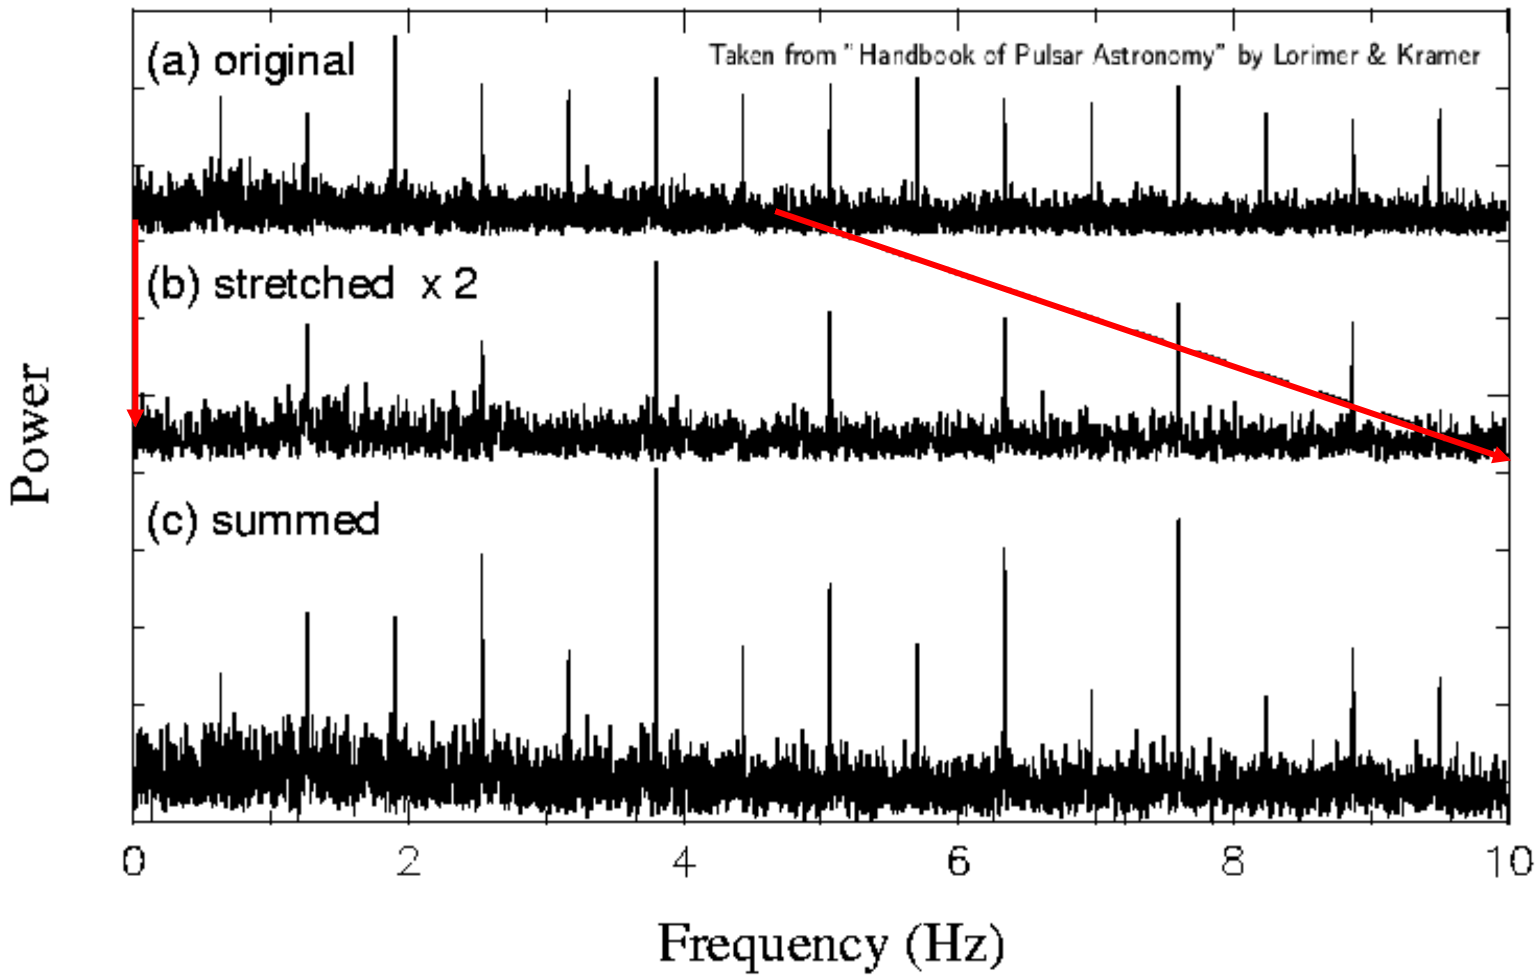

The number of harmonics to sum to get the best result depends on the duty cycle of the pulsar.

For broad profile pulsars, and for millisecond pulsars you're actually decreasing your sensitivity by adding harmonics! There is barely any power in them, and summing increases the noise level.

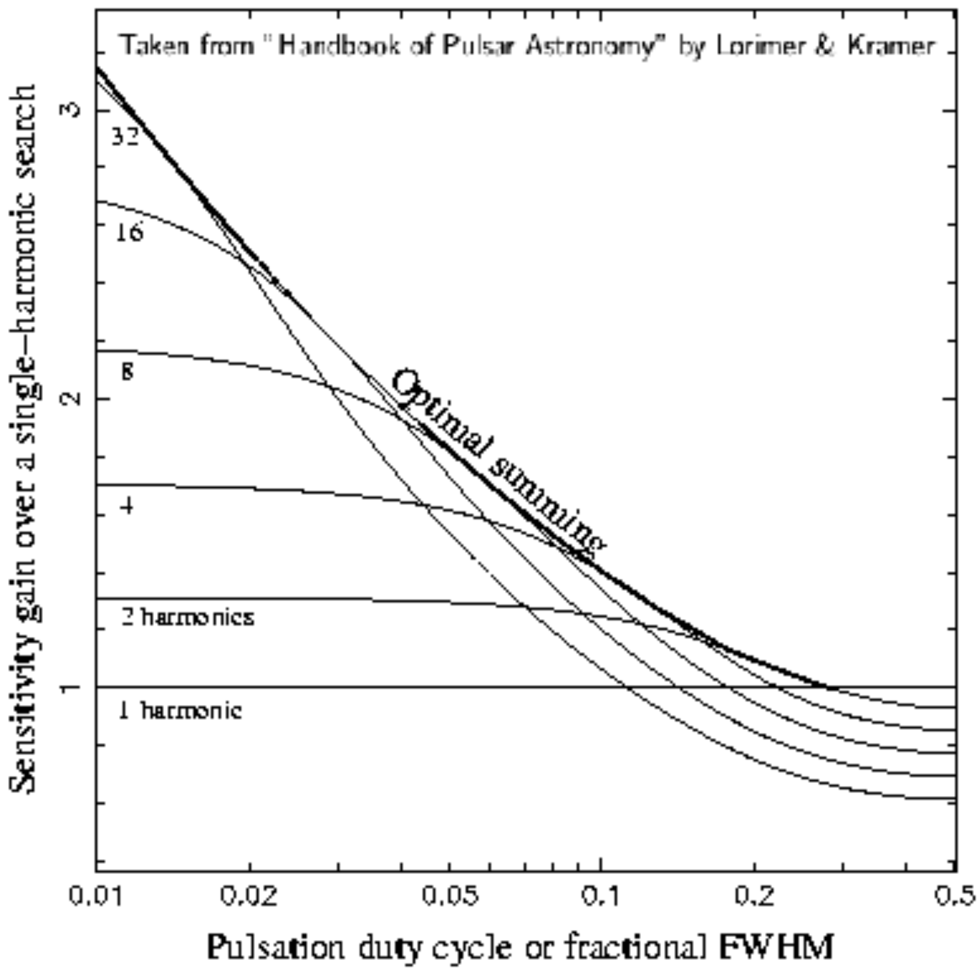

But usually the duty cycles are lower than 0.1, and in such case every little bit of gain helps.

Then, when you have a power spectrum of the dedispersed timeseries, cleared of the RFI, summed over harmonics - you may start to search for potential pulsar candidates.

And you have to use the data from several (or all) trial DMs at each time - since the pulsar candidate should show the S/N
varying with the DM.

So, you basically look for the candidates that show a maximum S/N at some non-zero DM.

Then you fold the timeseries with the trial period and check, if it looks as a pulsar.

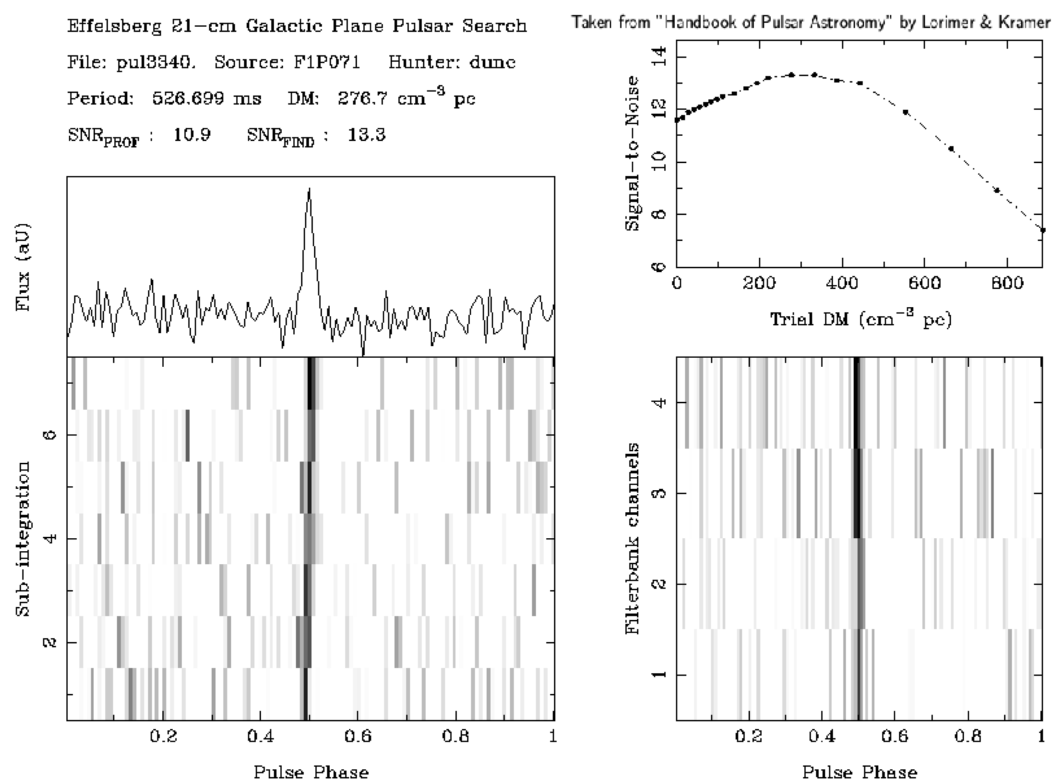

How to calculate your search sensitivity?

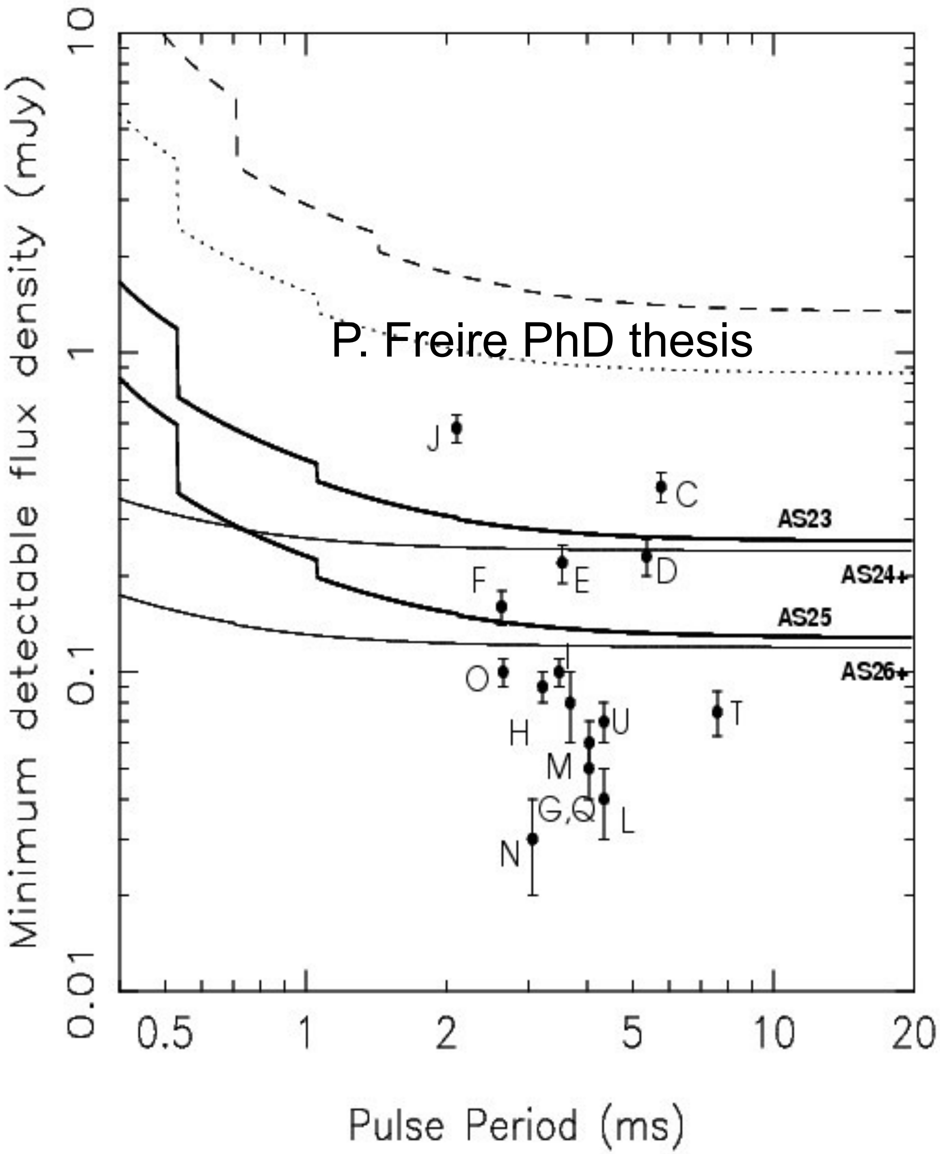

Minimum detectable flux depends on many factors:

$$S_{\rm min} = \frac{(S/N_{\rm min})\beta T_{\rm sys}}{G\sqrt{n_p t_{\rm obs}\Delta f}}\sqrt{\frac{\delta}{1-\delta}},$$

You have the control over some of them - especially if you are building your own backend.

Speaking about the binary puslars - in most cases, especially for tight binaries (which are of course more interesting) the orbital movement may change the apparent pulsar period during your observation. Therefore it will be harder to find sich pulsar in all that noise.

One has to find a way aroudn that...

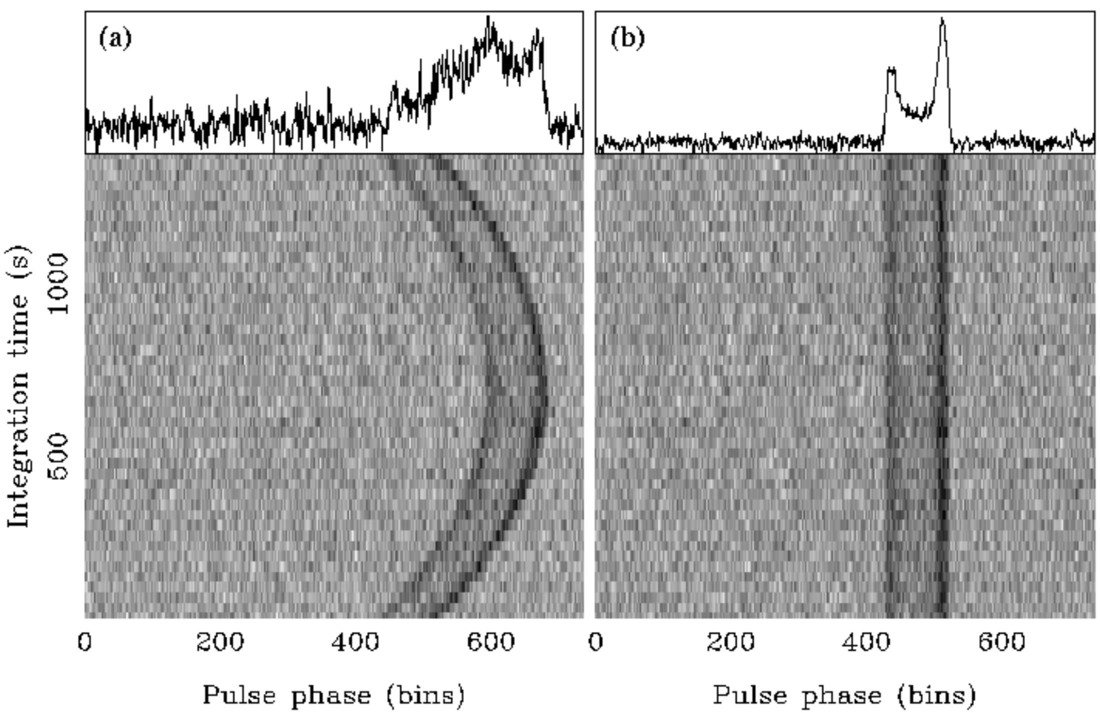

Taken from "Handbook of Pulsar Astronomy" by Lorimer & Kramer

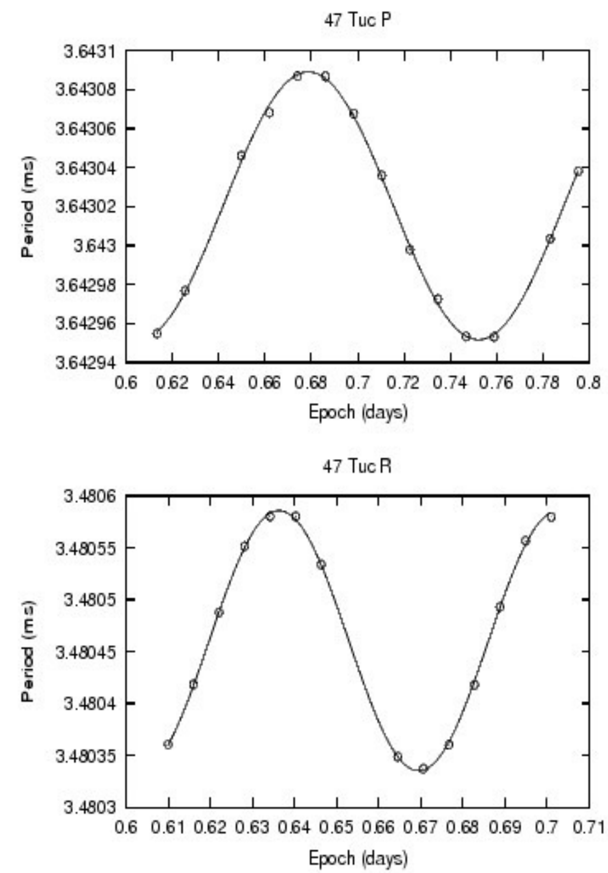

P. Freire PhD thesis

So basically, for millisecond binary pulsar you have to search not the usual 2-dimensional *P - DM* plane, but rather a 3-dimensional *P -DM -* $\dot{P}$ space!

And since you dont know beforehand what orbit to expect - it is a huge task.

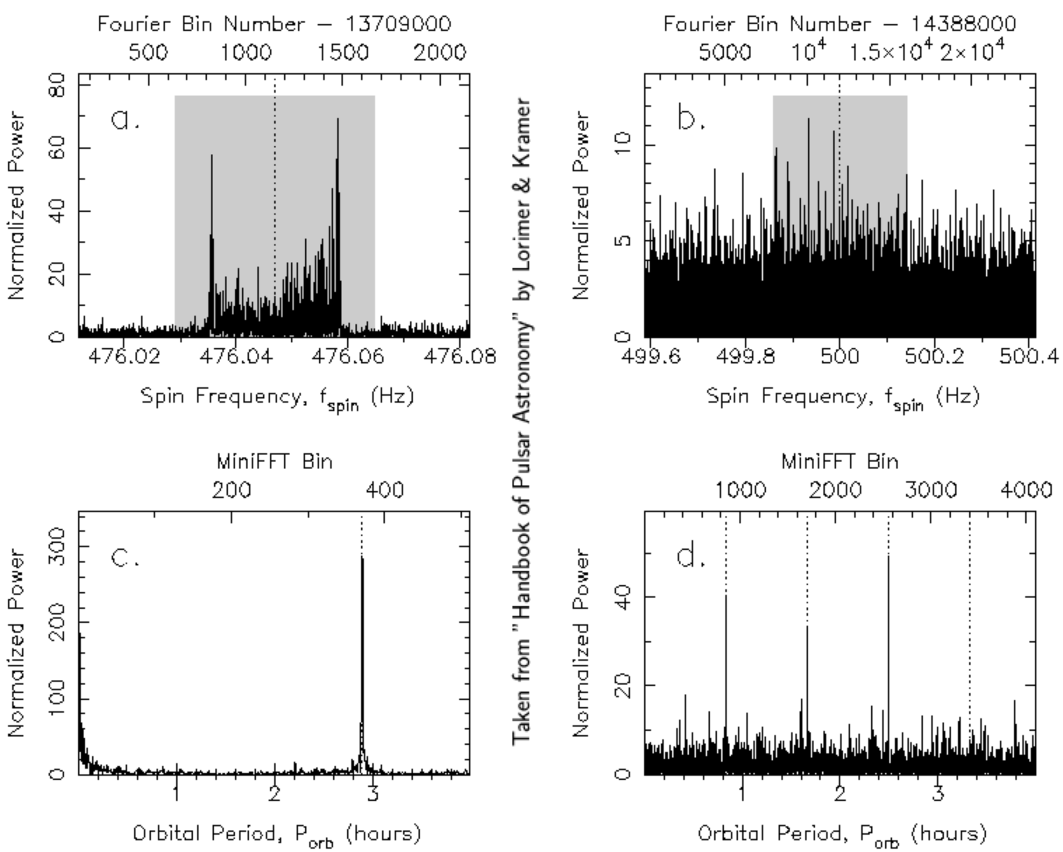

Mostly used is the padding method - you add/delete some bins from your timeseries to simulate the apparent period changes.

Other methods are Fourier based - such as the mini-FFT, performed on the fourier spectrum.

•A pulsar in a binary system experiences **orbital acceleration**

$$f_{\text{obs}}(t) = f_0 \left(1 + \frac{at}{c}\right)$$

•This causes **frequency drift during observation**, spreading power across Fourier bins.
•The acceleration search **re-corrects this drift** by trialing different accelerations.