
Gravitational Waves from Spinning Neutron Stars

in *Neutron Stars and Pulsars*, ed. W. Becker, Springer-Verlag 2009, p 651
The original publication is available at www.springerlink.com

LIGO-P060039-v3

Reinhard Prix¹ for the LIGO Scientific Collaboration

Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut,
Callinstr. 38, 30167 Hannover, Germany
Reinhard.Prix@aei.mpg.de

Summary. This paper gives an introductory overview of the search for continuous gravitational waves from spinning neutron stars. We review the current theoretical understanding of possible emission mechanisms and the expected strength of signals received on Earth. Given the substantial uncertainties involved in neutron star physics, these searches — if sufficiently sensitive — can provide important new information and constraints about neutron star physics even in the absence of detection. We describe the challenging detection problem of such extremely weak signals and introduce some of the basic data analysis concepts and methods used. We conclude by summarizing the current status of the sensitivities and upper limits achieved so far. These results suggest that the search for spinning neutron stars using LIGO-I is already entering a regime of “astrophysical relevance”. Developments in the near future make a detection of this type of source increasingly plausible.

1 Introduction

Gravitational waves (GWs), i.e. small deformations of spacetime traveling at the speed of light, are a fundamental consequence of Einstein’s General Theory of Relativity. There has been no direct observation of GWs so far, although first indirect evidence was found in the observed inspiral of the binary pulsar PSR 1913+16, which agreed to within 1% with the predictions of general relativity [88, 75]. Similar measurements on the recently discovered “double pulsar” system have allowed to improve these experimental tests of General Relativity to the level of 0.05%. GWs are purely transverse waves, characterized by two polarization states (denoted as ‘+’ and ‘×’, respectively). These two polarization bases differ by a rotation of 45 degrees around the propagation axis, corresponding to the quadrupolar (spin-2) nature of the gravitational field. In contrast, the two polarization bases of electromagnetic waves

differ by a rotation of 90 degrees, reflecting the dipolar (spin-1) nature of the electromagnetic field.

Any likely sources of detectable GWs will be at astrophysical distances, thus the signals reaching Earth have very small amplitudes and are nearly plane waves. A linearized version of general relativity (e.g. see [57]) can therefore be used to describe GWs in terms of a small metric perturbation $h_{\mu\nu}$, i.e. one can write the metric as $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$, where $|h_{\mu\nu}| \ll 1$ is the gravitational wave and $\eta_{\mu\nu}$ is the Minkowski metric of the unperturbed flat spacetime. One can then show that the Einstein field equations in vacuum reduce to the familiar wave-equation for a perturbation $h_{\mu\nu}$ propagating (at the speed of light) through flat spacetime $\eta_{\mu\nu}$, i.e.

$$\square h_{\mu\nu}^{\text{TT}} = \eta^{\sigma\rho} \partial_\sigma \partial_\rho h_{\mu\nu}^{\text{TT}} = 0, \quad (1)$$

where $h_{\mu\nu}^{\text{TT}}$ is the tensor $h_{\mu\nu}$ expressed in the transverse-traceless (TT) gauge. This is a choice of coordinates, $\{t, x, y, z\}$ say, corresponding to an inertial (Lorentz) frame in the unperturbed flat background, which makes explicit that the perturbation is *transverse*, i.e. orthogonal to the direction of propagation, and *trace-less*, namely the perturbation does not “compress” or “expand” elements of spacetime, but induces a (volume-preserving) “strain” only. In this gauge a plane gravitational wave propagating along the z-axis can be written as

$$h_{\mu\nu}^{\text{TT}}(t, z) = h_{\mu\nu}^{\text{TT}}(t - z/c) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (2)$$

where c is the speed of light and $h_{+, \times}(t - z/c)$ are the two polarizations of the wave. The effect of such a GW on two freely falling test-masses is a time-dependent change δl in their spatial distance l , which can be monitored using laser interferometry. This is the principle behind interferometric GW detectors such as the currently operating LIGO-I, GEO600, Virgo and TAMA300 (see Sec. 3), or future detectors such as Advanced LIGO or the space-based LISA. The measured strain $h(t)$ of such a detector is defined as the relative length-change between the two interferometer arms l_1 and l_2 (usually at right angles to each other), namely $h(t) \equiv \delta l_1/l_1 - \delta l_2/l_2$. In the long-wavelength approximation, where the reduced wavelength of the GW is assumed to be much longer than the arm length of the detector, i.e. $\lambda/2\pi \gg l_{1,2}$, it can be shown [69, 77, 47] that the measured scalar strain $h(t)$ due to the effect of a GW $h_{\mu\nu}^{\text{TT}}$ is given by

$$\begin{aligned} h(t) &= \frac{1}{2} \left(e_1^i e_1^j - e_2^i e_2^j \right) h_{ij}^{\text{TT}} \\ &= F_+(t) h_+(t) + F_\times(t) h_\times(t), \end{aligned} \quad (3)$$

where e_1^i and e_2^i are spatial unit-vectors along the first and the second interferometer arm, respectively. Note that the long-wavelength approximation

is generally sufficient for ground-based detectors with typical arm-lengths of $L \lesssim 4$ km and GW frequencies of up to \lesssim kHz. However, in the case of the space-based LISA detector, with armlengths $\sim 5 \times 10^6$ km and GW frequencies of order $\sim 10^{-2}$ Hz, the phase of the GW changes substantially during the light-travel time of photons between the space-craft. Therefore LISA requires a more detailed modeling of its response to a passing GW instead of the simple response tensor (3), which is achieved by “time-delay interferometry” or TDI (e.g. see [78]).

The antenna-pattern (or beam-pattern) functions $F_{+,\times}(t; \psi, \mathbf{n}) \in [-1, 1]$ describe the sensitivity of the detector to the ‘+’ and ‘×’ polarization, respectively, which depend on the direction \mathbf{n} to the source (where $\mathbf{n}^2 = 1$), the polarization angle ψ of the wave, and the orientation of the detector. For ground-based detectors, $F_{+,\times}$ are periodic functions over one sidereal day, due to the rotation of the Earth (while for LISA the period would be one year). Explicit expressions for the antenna-pattern functions can be found in [23] and [49]. GW detectors are fundamentally different from “telescopes” for electromagnetic radiation: they are practically omni-directional due to the wide quadrupolar antenna-pattern (3), and they measure a single scalar function of time, $h(t)$. In this sense they are more closely analogous to acoustic microphones rather than telescopes.

The emission of GWs is generally well-described by the quadrupole formula, namely

$$h_{jk}^{\text{TT}}(t) = \frac{2G}{c^4} \frac{1}{r} \left[\ddot{\mathcal{J}}_{jk}(t - r/c) \right]^{\text{TT}}, \quad (4)$$

where G is Newton’s gravitational constant, c is the speed of light, and r is the distance to the source and r/c accounts for the time delay. Dots denote time derivatives and \mathcal{J}_{jk} is the mass-quadrupole moment of the source. The transverse-traceless operator ‘TT’ indicates the projection orthogonal to the direction of propagation and the removal of the trace. This expression was first derived by Einstein in 1916, using the assumption of weak internal gravity of the source, but it was later shown to be valid even if one only requires that the source is small compared to the reduced wavelength $\lambda/2\pi$ of the emitted waves [77, 76]. The mass-quadrupole moment \mathcal{J}_{jk} of the source is defined as the coefficient of the $1/r^3$ term in an expansion in powers of $1/r$ of the Newtonian gravitational potential (far from the source). For sources with weak internal gravity, this can be directly expressed as

$$\mathcal{J}_{jk} = \int \rho(\mathbf{x}) \left[x_j x_k - \frac{1}{3} r^2 \delta_{jk} \right] d^3x, \quad (5)$$

where $\rho(\mathbf{x})$ is the mass density. This is simply the Newtonian moment of inertia with the trace removed. The energy emission rate L_{GW} in GWs can be expressed in the quadrupole formalism as

$$L_{\text{GW}} = \frac{1}{5} \frac{G}{c^5} \left\langle \ddot{\mathcal{J}}_{jk} \ddot{\mathcal{J}}^{jk} \right\rangle, \quad (6)$$

where $\langle \dots \rangle$ denotes the time average over several periods. The quadrupole formalism shows that time-varying mass-distributions generally emit GWs. Let us now specialize to the case of a star with mass M and radius R , rotating at a frequency ν . We see from (4) that a perfectly axisymmetric star rotating around its symmetry axis will *not* emit GWs, as its quadrupole moment (5) is constant in time. Let us characterize the deviation from axisymmetry by a dimensionless number ϵ , i.e. let ϵI be the non-axisymmetric part of the moment of inertia I . Then the order of magnitude of the GW luminosity (6) can be estimated as

$$L_{\text{GW}} \sim \frac{G}{c^5} \epsilon^2 I_{zz}^2 \nu^6, \quad (7)$$

where $I_{zz} \sim M R^2$ is the moment of inertia along the rotation axis. Numerically $G/c^5 \sim 3 \times 10^{-53}$ s/J, which shows that terrestrial experiments could realistically only produce infinitesimal amounts of GW radiation: if we constructed a “dumbbell” of two masses of a few tons separated by a few meters, i.e. an $\epsilon \sim \mathcal{O}(1)$ and a moment of inertia $I_{zz} \sim 10^5 \text{ kg m}^2$, and if we let it spin or vibrate at rates $\nu \sim 10^3 \text{ Hz}$, it would still radiate only about $L_{\text{GW}} \sim 10^{-25}$ W in gravitational waves. On the other hand, if we rewrite this expression in terms of the Schwarzschild radius $R_s = 2GM/c^2$ and the rotation velocity $V = 2\pi R\nu$ at the surface of the star [24], we find

$$L_{\text{GW}} \sim \frac{c^5}{G} \epsilon^2 \left(\frac{R_s}{R} \right)^2 \left(\frac{V}{c} \right)^6. \quad (8)$$

This shows that *compact objects* (i.e. $R_s \sim R$) in rapid rotation (i.e. $V \sim c$), such as spinning neutron stars, can emit enormous GW luminosities even for small ϵ , as $c^5/G \sim 10^{52}$ W. However, spacetime is a very “stiff” medium, i.e. large amounts of energy are carried by GWs of small amplitude. This can be seen from a similar estimate of the corresponding strain amplitude (4), namely

$$h \sim 10^2 \frac{G}{c^4} \frac{\epsilon I_{zz} \nu^2}{d} \sim 3 \times 10^{-25} \left(\frac{\epsilon}{10^{-6}} \right) \left(\frac{I_{zz}}{10^{38} \text{ kg m}^2} \right) \left(\frac{\nu}{100 \text{ Hz}} \right)^2 \left(\frac{100 \text{ pc}}{d} \right), \quad (9)$$

where the fiducial values correspond to a neutron star with typical moment of inertia, a relatively strong deviation from axisymmetry (cf. next section), at a distance of the order of the closest known neutron star, and spinning in the millisecond regime. Even in this very optimistic case, the amplitude is about two orders of magnitude below the noise level (referring to 1 s of data) of LIGO-I, which is currently the most sensitive detector and will be discussed in more detail in Sec. 3. The data-analysis challenge in the search for continuous GWs from neutron stars therefore consists of finding extremely weak, nearly periodic signals buried deep in the noise of GW detectors.

2 Continuous gravitational waves from neutron stars

In this section we give a brief overview of the current theoretical understanding of various physical mechanisms that could operate in neutron stars to produce interesting levels of GW emission (see also [61] for a recent review). As the signal-strength is generally expected to be very weak (see previous section), long integration times of the order of $T \gtrsim$ several days to years will be required in order for the signal to be detectable in the noise (see Sec. 3).

Here we focus on the class of *continuous GWs*, which are long-lasting, quasi-monochromatic GWs with slowly varying intrinsic frequency. The difficulties and analysis methods for these signals differ in many respects from other types of GW signals, such as bursts (e.g. from supernovae), “chirping” signals from the binary coalescence of compact objects, and the fossil stochastic background of GWs from the big bang. Note that neutron stars might also be interesting sources of burst-like GW emission [33], e.g. from oscillations excited by a glitch which would be damped very quickly. Another candidate for such burst-like neutron star oscillations could be crustal torsional modes, which have recently been suggested [87] as a possible explanation for the high frequency oscillations observed in giant flares from Soft Gamma Repeaters. Although such GW “asteroseismology” could potentially be extremely valuable for neutron-star astrophysics (e.g. see [14]), a discussion of this type of sources is not within the scope of the present paper, as the corresponding detection problem is rather different in nature from the continuous-wave type of sources considered here.

2.1 Emission mechanisms for continuous gravitational waves

There are three types of mechanisms usually considered for the emission of continuous GWs from spinning neutron stars in the frequency band of current ground-based detectors (i.e. ~ 20 Hz–2kHz): (i) non-axisymmetric distortions of the neutron star, (ii) unstable oscillation modes in the fluid part of the star (e.g. r-modes), and (iii) free precession.

Non-axisymmetric distortions

Non-axisymmetric distortions, or “mountains” on the neutron star, cannot persist in perfect fluid stars, but in realistic neutron stars such deformations can be supported either by elastic stresses in the crust or by magnetic fields. The deformation is often expressed in terms of the *equatorial* ellipticity:

$$\epsilon \equiv \frac{I_{xx} - I_{yy}}{I_{zz}}, \quad (10)$$

where I_{jj} are the three principal moments of inertia. This ellipticity is not to be confused with the centrifugal bulge, which is axisymmetric and therefore

does not radiate GWs (see Sec. 1). A spinning non-axisymmetric neutron star at distance d , rotating with frequency ν around the I_{zz} axis emits monochromatic GWs at a frequency $f = 2\nu$ of amplitude

$$h_0 = \frac{16\pi^2 G}{c^4} \frac{I_{zz} \nu^2}{d} \epsilon, \quad (11)$$

e.g. see [49], and Eq. (9). The definition of the strain-amplitude h_0 refers to a GW from an optimally oriented source with respect to the detector, as will become clearer in Sec. 4.1.

Unfortunately, both the maximum possible as well as the most likely values for the non-axisymmetric deformation ϵ of neutron stars are highly uncertain. Various estimates of the maximum ellipticity ϵ sustainable by neutron stars have been given in the literature: the maximal deformation supported by the rigidity of the crust has been estimated [81] as

$$\epsilon_{\max} \approx 5 \times 10^{-7} \left(\frac{\sigma}{10^{-2}} \right), \quad (12)$$

where σ is the breaking strain of the solid crust. The numerical coefficient in this equation is small mainly because the shear modulus of the inner crust is small compared to the pressure. The fiducial breaking strain of 10^{-2} in Eq. (12) corresponds approximately to the best terrestrial alloys. However, σ could be as high as 10^{-1} for a perfect crystal with no defects [51], or several orders of magnitude smaller for an amorphous solid or a crystal with many defects. Some exotic alternatives to standard neutron star models contain solid cores, which could support considerably larger ellipticities. The most speculative models for solid strange-quark stars can reach ellipticities of up to $\epsilon_{\max} \approx 4 \times 10^{-4} (\sigma/10^{-2})$, mostly due to their higher shear modulus [60].

In addition to the problem of the maximum ellipticity that can be supported, the more relevant question is what actual deformations are likely to exist in real neutron stars. There are several mechanisms by which the spin-down of a neutron star could cause stresses in the crust to build up to the point of breaking [see 66, 67, 28], usually considered in the context of glitch models. It is unclear, however, how long it would take to return to a smooth axisymmetric mass-distribution after such a crust quake, which could possibly leave long-lived distortions of the crust.

Another possibility are accreting neutron stars in binary systems, which have a natural way of reaching and maintaining large crust deformations: the accretion flow, guided by the neutron star's magnetic field, naturally produces "hot spots" on the surface, which can lead to "hills" in hotter areas, and the ellipticity might thereby build up to the maximum value [22]. The accreted material could also be held up in mountains by the magnetic field itself: the accreted matter has a very high electric conductivity, thus crosses field lines relatively slowly and could pile up in mountains larger than those supportable by elasticity alone [56, 64]. Depending on the field configuration, accretion

rate, and temperature, the ellipticity from this mechanism could be up to a few times 10^{-6} even for ordinary neutron stars (cf. [61] for further discussion).

Furthermore, strong internal magnetic fields could directly produce non-axisymmetric deformations of neutron stars due to magnetic tension. A strong dipolar field that is not aligned with the rotation axis could lead to deformations of up to $\epsilon \lesssim 10^{-6}$ in the case of type-I superconducting cores [23]. These non-aligned deformations would generally result in GW emission at *both* the first and the second harmonic of the rotation rate, i.e. $f = \nu$ and $f = 2\nu$, similar to free precession. Another possibility, suggested by [35], is that the differential rotation present after the birth of a neutron star could “wind up” large toroidal magnetic fields. Dissipation then tends to drive the symmetry axis of a toroidal field toward the star’s equator, resulting in ellipticities of the order $\epsilon \sim 10^{-6}$ for toroidal magnetic fields of about $B_t \sim 10^{15}$ G.

Non-axisymmetric instabilities

At birth or during phases of accretion, rapidly rotating neutron stars can be subject to various non-axisymmetric instabilities, which would lead to GW emission (see [12, 73] for recent reviews). If a new-born neutron star has a sufficiently high rotation rate, i.e. if the ratio $T/|W|$ of the rotational kinetic energy T to the gravitational binding energy W exceeds a critical value (typically $T/|W| \gtrsim 0.24$), it will be subject to a dynamical instability driven by hydrodynamics and gravity. Differential rotation might substantially lower this critical value to $T/|W| \gtrsim 0.14$ [29]. Through the $l = 2$ mode, this bar-mode instability will deform the star into a bar shape, which would be a strong emitter of GWs, but likely to be very short-lived (on the order of a few rotations [20]). At lower rotation rates, the star can become unstable to *secular* non-axisymmetric instabilities, driven either by gravitational radiation or viscosity [e.g. 70, 68]. It is not clear, however, how effective and long-lived any of these mechanisms would be in producing detectable GWs.

Gravitational radiation generally drives a non-axisymmetric oscillation unstable if the mode is counter-rotating with respect to the rotating frame of the star, but co-rotating with the star in the frame of a distant inertial observer. This happens when the counter-rotation rate of the mode is lower than the rotation rate of the star, so the mode gets “dragged along” by the star. This secular instability is known as the Chandrasekhar-Friedman-Schutz (CFS) instability [31, 44]. For oscillation modes such as the fundamental mode (f-mode) and pressure-modes (p-modes), this instability would only set in at very high rotation rates close to the breakup-limit. Therefore the most promising candidate for observable GW emission from the CFS instability might be the family of r-modes, which are toroidal fluid oscillations where the Coriolis force of the rotating star is the dominant restoring mechanism (see [80] for an overview of stellar oscillations). Contrary to the polar f- and p-modes, these oscillations are generically unstable to the CFS instability at any finite rotation rate [11, 43].

Under astrophysically realistic conditions, however, their effective instability depends on a number of highly uncertain damping mechanisms and timescales (see [73] for a review). The r-mode instability has been proposed as a source of GWs (with frequency $f = 4\nu/3$) from newborn neutron stars [62] and from rapidly accreting neutron stars [22, 16]. However, the CFS instability of the r-modes in newborn neutron stars might not be a good candidate for detection because the emission will most likely be weak and would last only for a few months after the birth of the neutron stars in a supernova. Accreting neutron stars or quark stars could be a better prospect for detection [85, 13] because the emission may last several thousands of years [15, 12].

Free precession

The third major type of emission mechanisms for continuous GWs from spinning neutron stars is free precession, i.e. the “wobble” of a neutron star with a misaligned rotation axis with respect to its symmetry axis (defining the wobble angle θ_w). A large-angle wobble would typically [see 90, 50, 82] produce GWs of amplitude

$$h_0 \sim 10^{-27} \left(\frac{\theta_w}{0.1 \text{ rad}} \right) \left(\frac{1 \text{ kpc}}{d} \right) \left(\frac{\nu}{500 \text{ Hz}} \right)^2. \quad (13)$$

Generally, free precession results in emission at (approximately) the rotation rate ν and twice the rotation rate, i.e. $f = \nu + \nu_{\text{prec}}$ (where ν_{prec} is the precession frequency) and $f = 2\nu$. Free precession may be much longer lived ($\sim 10^5$ years) than previously thought [37], but the amplitude is still quite small, which might make this mechanism relevant only for second-generation interferometers such as Advanced LIGO.

2.2 Loudest expected signal from unknown isolated neutron stars

Current models of stellar evolution suggest that our Galaxy contains $\sim 10^9$ neutron stars, while about 10^5 are expected to be active pulsars. Up to now only ~ 1700 pulsars have been observed [see 54]; there are numerous reasons for this, including selection effects and because many produce only faint emission. Although there is great uncertainty in the physics of the GW emission mechanisms and the strength of individual sources, one can argue for a statistical upper limit on the expected strongest GW signals from the galactic population of neutron stars. The argument is almost independent of individual source physics and goes back to Blandford. An updated version is given in [7], here we only summarize the main points and the result. One can make the (optimistic) assumption that there exists a class of neutron stars that are born rapidly rotating, and spinning down due to GWs. Assuming they are distributed uniformly throughout the galactic disc, and have a constant birthrate, one can convert these assumptions into a distribution of neutron

stars with respect to GW strain and frequency. Using this, one can show that there would be a 50% chance that the strongest signal within the LIGO band ($\sim 50 \text{ Hz} - 2 \text{ kHz}$) has an amplitude of at least

$$h_0 \sim 4 \times 10^{-24}. \quad (14)$$

Obviously this is a purely statistical argument, and it is possible that the closest strong emitter is either far closer or further away than the typical distance expected from a random distribution of supernovae. Note also that this is an “upper limit” only in the sense that the assumptions that went into it are optimistic, and one would generally expect a lower amplitude if some of them were not true.

2.3 The spindown limit for known pulsars

The statistical argument of the previous section applies to a population of unknown neutron stars in the galaxy. A much more robust upper limit on h_0 can be derived for known pulsars with measured frequency ν , spindown $\dot{\nu}$ and distance d . For simplicity, let us consider the case of spinning neutron stars with a non-axisymmetric deformation ϵ , emitting GWs at frequency $f = 2\nu$. In this case the GW luminosity (6) can be written explicitly as

$$L_{\text{GW}} = \frac{1}{10} \frac{G}{c^5} (4\pi\nu)^6 I_{zz}^2 \epsilon^2. \quad (15)$$

If we assume that the GW emission is powered *only* by the rotational energy $E_{\text{rot}} = 2\pi^2\nu^2 I_{zz}$ of the spinning neutron star, then we have the inequality

$$L_{\text{GW}} \leq -\dot{E}_{\text{rot}} = -2\pi^2 \left(2I_{zz} \nu \dot{\nu} + \nu^2 \dot{I}_{zz} \right). \quad (16)$$

For constant moment of inertial, $\dot{I}_{zz} = 0$, this yields an upper limit on the quadrupolar deformation $\epsilon \leq \epsilon_{\text{sd}}$, namely

$$\epsilon_{\text{sd}} = \sqrt{\frac{5c^5}{2(4\pi)^4 G I_{zz}} \frac{|\dot{\nu}|}{\nu^5}}. \quad (17)$$

Substituting this into Eq. (11), we obtain the following upper limit on the GW amplitude

$$h_0 \leq h_{\text{sd}} = \frac{1}{d} \sqrt{\frac{5G I_{zz}}{2c^3} \frac{|\dot{\nu}|}{\nu}}. \quad (18)$$

2.4 Maximum expected signal from accreting neutron stars

The statistical upper limit (14) on h_0 is not applicable to accreting neutron stars, since energy conservation plays a crucial role in the argument. However, if accretion replenishes the star’s angular momentum, a different argument can

be made, independent of the details of the emission mechanism. In this case h_0^{\max} is set by the X-ray luminosity of the brightest X-ray source. The basic idea is that if the angular momentum lost in GWs is replenished by accretion, then the strongest GW emitters are those accreting at the highest rate, such as low-mass X-ray binaries (LMXBs). The accreted gas hitting the surface of the neutron star is heated to $\sim 10^8$ K and emits X-rays. As noted several times [63, 84, 22], if one assumes the spindown torque from GW emission to be in equilibrium with the accretion torque, then h_0 is directly related to the observed X-ray flux F_x , namely

$$h_0 \approx 5 \times 10^{-27} \left(\frac{300 \text{ Hz}}{\nu} \right)^{1/2} \left(\frac{F_x}{10^{-8} \text{ erg cm}^{-2} \text{ s}^{-1}} \right)^{1/2}. \quad (19)$$

It is interesting to note that this expression does not explicitly depend on the distance d of the system, which is implicitly contained in the observed X-ray flux F_x on Earth (both the GW and X-ray flux decrease with $1/d^2$). This theoretical argument is supported by the observation that the frequencies of most LMXBs seem to cluster in a fairly narrow range of $270 \text{ Hz} \lesssim \nu \lesssim 620 \text{ Hz}$ [see 30]. Since most neutron stars will have accreted enough matter to spin them up to near their theoretical maximum spin frequencies, estimated at $\nu_{\max} \sim 1400 \text{ Hz}$ (e.g. see [34]), the observed spin distribution is hard to explain without a competing mechanism to counter the spin-up. Since the GW torque scales as $\propto \nu^5$, gravitational radiation seems like a natural explanation for the rather narrow clustering of observed frequencies. If this argument holds, then the accreting neutron star brightest in X-rays (namely Sco X-1) should also be the strongest source of GWs. Using the known X-ray flux of Sco X-1, its expected GW emission under these assumptions would have an amplitude of

$$h_0 \approx 3 \times 10^{-26} \left(\frac{540 \text{ Hz}}{f} \right)^{1/2}. \quad (20)$$

The frequency f is unfortunately not well constrained from observations, but is generally assumed to be of the order of several hundred Hz. This signal could in principle be detectable by second-generation interferometers such as Advanced LIGO (cf. Fig. 4).

3 Detectors of Gravitational Waves

Starting from the pioneering efforts of Joseph Weber in the early 1960's, the first GW detectors were based on the principle of monitoring the oscillations of massive resonant metal bars, the bar detectors. A GW at or near the resonance frequency of the bar would excite this oscillation mode. These designs have been successively improved over time, and today there are still a number of bar detectors operating, including ALLEGRO in Louisiana, EXPLORER

at CERN and NAUTILUS in Rome. During the past decade, however, several scientific collaborations have constructed large-scale interferometric GW detectors. These include the Laser Interferometer Gravitational Wave Observatory (LIGO), consisting of three interferometers, built by a Caltech-MIT collaboration [9]; the GEO 600 detector built by a British-German collaboration [89]; the Virgo detector built by an Italian-French collaboration [27]; and the Japanese TAMA 300 detector in Tokyo [79]. In all of these detectors, the relative displacement of suspended test masses is sensed interferometrically, as illustrated schematically in Fig. 1. The interferometer is tuned in such a




Fig. 1. Schematic layout of an interferometric gravitational-wave detector such as LIGO, with power recycling and Fabry-Perot arms. The quadrupolar strain deformation-field above the detector indicates an incident GW $h_{\mu\nu}^{TT}$ of optimal orientation (from [1]).

way that ideally no light would arrive at the photodetector in the absence of a GW. Therefore most of the light will be reflected back to the laser, where it is returned to the interferometer using a power recycling mirror. This increases the power of the light in the interferometer arms, which reduces “shot noise”, i.e. the statistical fluctuations of the laser light due to the quantum nature of photons.

In addition to shot noise, which is the dominant noise contribution at high frequencies, the strain measurement is affected by a large variety of noise sources. At low frequencies the dominant contributions are seismic noise and gravity-gradient noise. See [77, 1] for more detailed discussions about the functioning of GW detectors and the problems related to minimizing noise contributions affecting the strain measurements.

In the following we will focus mainly on the detectors run within the LIGO Scientific Collaboration (LSC), namely LIGO and GEO 600. The GEO 600 detector (G1) comprises a 4-beam Michelson delay line system of arm length 600 m. LIGO consists of three power-recycled Michelson interferometers with resonant Fabry-Perot cavity arms, installed at two sites: the Livingston site (Louisiana) contains one interferometer of 4 km arm length (referred to as L1), while the Hanford site (in Washington state) houses two interferometers, one of 4 km and one of 2 km arm length (called H1 and H2, respectively). In all four instruments (H1, H2, L1, G1), the beam splitters, recycling mirrors and test masses are hung as pendulums from multilayer seismic isolation filters to isolate them from local forces (see Fig. 1). The masses and beam paths are housed in high vacuum enclosures to avoid optical scintillation and acoustic interference.

3.1 LIGO/GEO600 Sensitivities and Scientific Runs

LIGO and GEO600 have so far completed five science-mode data-taking runs (denoted S1-S5), see Table 1. The “duty cycle” in Table 1 denotes the fraction

Table 1. Summary of LIGO/GEO600 scientific runs. [†]Partial participation: Nov 5, 03 – Nov 12, 03 and Dec 30, 03 – Jan 13, 04, [‡]started Nov 14, 05.

Run	Start – End	best sensitivity $\sqrt{S_n}$ [Hz ^{-1/2}]	duty cycles [%]			
			H1	L1	H2	G1
S1	Aug 23 – Sept 9, 2002	2×10^{-21}	57.6	41.7	73.1	98.5
S2	Feb 14 – Apr 14, 2003	2×10^{-22}	73.5	36.9	57.8	–
S3	Oct 31, 2003 – Jan 9, 2004	5×10^{-23}	69.3	21.8	63.4	96.9 [†]
S4	Feb 22 - Mar 23, 2005	4×10^{-23}	80.5	74.5	81.4	96.6
S5	Nov 4, 2005 - Oct 1, 2007	2×10^{-23}	71.0	59.1 [‡]	78.2	80.4

of the run time where the detector was “in lock” and was taking science data. Due to seismic noise, equipment failures and alignment drifts, the duty cycle is generally less than 100%. Livingston (L1) had particularly low duty-cycles during S1 to S3 due to low-frequency noise caused by logging activity in a nearby forest. This problem has been largely overcome by installing an improved (active) seismic isolation before the start of S4, which resulted in dramatic improvements in the L1 duty cycle, as seen in Table 1. GEO 600 has had the best duty cycles, but also lower sensitivity than LIGO, see Fig. 2. As seen in this figure, the LIGO detectors have reached their design sensitivity (except at low frequencies $\lesssim 50$ Hz) during the S5 run, which took one year of coincident data (the run lasted for about 2 years). After S5, LIGO-I will undergo some enhancements (known as “Enhanced LIGO”) that will improve sensitivity by up to a factor of two, and will be used for another long science run (S6). After S6, work will start on the next-generation Advanced LIGO




Fig. 2. *Left figure:* Successive best LIGO sensitivities achieved over the science runs S1 - S5. *Right figure:* LIGO and GEO600 sensitivities during early S5 (June 2006).

detector, which will be installed on the same sites as LIGO-I, and which is planned to start taking data within a decade.

4 Data analysis of continuous gravitational waves

As discussed in Secs. 1 and 2, continuous GWs reaching Earth are expected to be exceedingly weak, even compared to the sensitivities of the current generation of detectors. In order to be able to dig such signals out of the noise, one typically has to integrate for several days up to months by matching the data with a target signal (“template”) of given parameters. This is the basic concept of *matched filtering*, which is the optimal method in a statistical sense (made more precise later). See also [48] for a recent review of the detection problem of continuous GWs.

As noted earlier, GW detectors are practically omni-directional. Due to the long integration time, however, it turns out that continuous GWs are in fact extremely well-localized, not only in frequency but also in sky-position. A mismatch in frequency between the true signal and a template results in a phase mismatch that is growing with observation time, and thereby rapidly degrading the output of the (mis-)matched filter. Similarly, the Doppler effect from the daily rotation and orbital motion of the Earth modulates signals in a way that depends sensitively on the direction from which they are coming. The templates therefore need to be sky-position specific, and the required precision increases with observation time.

In the case of *wide-parameter searches* for unknown sources, this makes it very expensive in terms of computing cost to increase the observation time, as it requires a much finer search in the parameter space of possible signals. There is also a purely statistical effect limiting the sensitivity of such searches: the more trials (i.e. targeting of different points in parameter space) one per-

forms, the more “false alarm” candidates are expected to cross a given detection threshold due to noise fluctuations alone. Therefore, a higher detection threshold is required, which reduces the sensitivity.

On the other hand, fully *targeted searches* for GW sources with known parameters (such as pulsars with known sky-position and frequency evolution) are not affected by these difficulties and can attain the best possible sensitivity by coherently integrating over all the available data.

4.1 The general form of the signal

A general, elliptically polarized GW can be written in the reference frame of the source as

$$h_+(\tau) = A_+ \cos \Phi(\tau), \quad h_\times(\tau) = A_\times \sin \Phi(\tau), \quad (21)$$

where $h_{+,\times}$ are the two polarization states of $h_{\mu\nu}^{\text{TT}}$ given in (2), using coordinates aligned with the principal polarization axes of the coherent GW. Assuming a quasi-monochromatic signal with slowly-varying intrinsic frequency $f(\tau)$, the signal phase $\Phi(\tau)$ can be Taylor-expanded as

$$\Phi(\tau) = \phi_0 + \phi(\tau), \quad \phi(\tau) = 2\pi \sum_{k=0}^s \frac{f^{(k)}(\tau_{\text{ref}})}{(k+1)!} \Delta\tau^{k+1}, \quad (22)$$

where $\Delta\tau \equiv \tau - \tau_{\text{ref}}$, and τ_{ref} is the reference time at which the initial phase ϕ_0 and the $s+1$ spin parameters $f^{(k)} \equiv d^k f(\tau)/d\tau^k$ are defined.

Let $\mathbf{n} = (\cos \delta \cos \alpha, \cos \delta \sin \alpha, \sin \delta)$ be the unit vector pointing to the source, expressed in equatorial coordinates using the standard celestial angles α (right ascension) and δ (declination). The wave-frame is then completely determined by \mathbf{n} and the “polarization angle” ψ , which describes the orientation of the polarization axes with respect to the equatorial-coordinate system. Following the conventions of [23], ψ can be defined as the angle between the direction $\mathbf{n} \times \mathbf{Z}$ and the \mathbf{x} -axis of the TT wave-frame (corresponding to the ‘+’ polarization), where $\mathbf{Z} = (0, 0, 1)$ is the unit-vector pointing to the celestial north pole. As discussed in [49], the dependency of the antenna-pattern functions $F_{+,\times}$ on the wave-frame orientation $\{\mathbf{n}, \psi\}$ can be separated as

$$\begin{aligned} F_+(t; \mathbf{n}, \psi) &= a(t; \mathbf{n}) \cos 2\psi + b(t; \mathbf{n}) \sin 2\psi, \\ F_\times(t; \mathbf{n}, \psi) &= b(t; \mathbf{n}) \cos 2\psi - a(t; \mathbf{n}) \sin 2\psi, \end{aligned} \quad (23)$$

where the expressions for the (detector-dependent) functions $a(t; \mathbf{n})$, $b(t; \mathbf{n})$ are given in Eqs. (12,13) of [49].

In the detector frame the signal amplitude is modulated by the rotating antenna pattern $F_{+,\times}(t)$, as seen in Eq. (3). More importantly, the signal is also Doppler-modulated by the relative motion of the detector with respect to the source. This can be expressed as a relation between the detector arrival

time t of a wave-front that left the source at time $\tau(t)$. Let us consider the most general case of a neutron star in a binary system with orbital parameters \mathbf{b} (including orbital period, projected semi-major axis, ellipticity etc) at a sky-position \mathbf{n} . The timing relation $\tau(t)$ can then be written as

$$\tau(t; \mathbf{n}, \mathbf{b}) = t + \frac{\mathbf{r}(t) \cdot \mathbf{n}}{c} - \frac{d}{c} + \Delta_{\text{bin}}(t; \mathbf{b}), \quad (24)$$

where $\mathbf{r}(t)$ is the vector from the solar-system barycenter (SSB) to the detector location, and $\mathbf{r} \cdot \mathbf{n}/c$ term is known as the Roemer-delay. For simplicity we have neglected relativistic corrections in the SSB such as the Shapiro and Einstein delays, see [75] for details. $\Delta_{\text{bin}}(t; \mathbf{b})$ is the time delay between the neutron star and its binary-system barycenter (BSB), and d is the distance between the BSB and the SSB, which for simplicity is assumed constant and can therefore be neglected. For isolated neutron stars we would obviously have $\Delta_{\text{bin}} = 0$. Inserting the timing relation (24) into the phase (22) in the source frame, we see that the phase evolution $\phi(t)$ in the detector frame has the general form

$$\phi(t; \boldsymbol{\lambda}) = \phi(\tau(t; \boldsymbol{\lambda})), \quad \text{where } \boldsymbol{\lambda} \equiv \{f^{(k)}, \mathbf{n}, \mathbf{b}\}, \quad (25)$$

which defines the set of *Doppler parameters* $\boldsymbol{\lambda}$. Putting all the pieces together, we can express the strain (3) in the detector as

$$\begin{aligned} h(t; \mathcal{A}, \boldsymbol{\lambda}) = & F_+(t; \mathbf{n}, \psi) A_+ \cos[\phi_0 + \phi(t; \boldsymbol{\lambda})] \\ & + F_\times(t; \mathbf{n}, \psi) A_\times \sin[\phi_0 + \phi(t; \boldsymbol{\lambda})], \end{aligned} \quad (26)$$

where we defined the four *amplitude parameters* $\mathcal{A} \equiv \{A_+, A_\times, \psi, \phi_0\}$. Using the form (23) of the antenna-pattern functions $F_{+, \times}$, it is now easy to see that the dependencies on the amplitude and Doppler parameters can be explicitly separated, namely

$$h(t; \mathcal{A}, \boldsymbol{\lambda}) = \sum_{\mu=1}^4 \mathcal{A}^\mu h_\mu(t; \boldsymbol{\lambda}), \quad (27)$$

in terms of four basis waveforms

$$\begin{aligned} h_1(t; \boldsymbol{\lambda}) &= a(t; \mathbf{n}) \cos \phi(t; \boldsymbol{\lambda}), & h_2(t; \boldsymbol{\lambda}) &= b(t; \mathbf{n}) \cos \phi(t; \boldsymbol{\lambda}), \\ h_3(t; \boldsymbol{\lambda}) &= a(t; \mathbf{n}) \sin \phi(t; \boldsymbol{\lambda}), & h_4(t; \boldsymbol{\lambda}) &= b(t; \mathbf{n}) \sin \phi(t; \boldsymbol{\lambda}), \end{aligned} \quad (28)$$

and the amplitude vector \mathcal{A}^μ , defined as

$$\begin{aligned} \mathcal{A}^1 &= A_+ \cos \phi_0 \cos 2\psi - A_\times \sin \phi_0 \sin 2\psi, \\ \mathcal{A}^2 &= A_+ \cos \phi_0 \sin 2\psi + A_\times \sin \phi_0 \cos 2\psi, \\ \mathcal{A}^3 &= -A_+ \sin \phi_0 \cos 2\psi - A_\times \cos \phi_0 \sin 2\psi, \\ \mathcal{A}^4 &= -A_+ \sin \phi_0 \sin 2\psi + A_\times \cos \phi_0 \cos 2\psi. \end{aligned} \quad (29)$$

Different emission mechanisms of continuous GWs result in different expressions for the amplitude parameters \mathcal{A} in terms of the source parameters, and in different relations between the rotation frequency ν of the neutron star and the GW frequency f in (22). Some interesting special cases are triaxial neutron stars rotating around a principal axis, in which case $f = 2\nu$. Free precession emits additional power at a frequency $f \approx \nu$, and r-mode oscillations emit near $f \approx 4\nu/3$. In order to simplify the following discussion, it is convenient to express the amplitudes $A_{+, \times}$ in terms of the source parameters of a non-precessing triaxial neutron star [e.g. 49], namely

$$A_+ = \frac{1}{2} h_0 (1 + \cos^2 \iota), \quad A_\times = h_0 \cos \iota, \quad (30)$$

where h_0 is the overall amplitude (11), and ι is the angle between the spin-axis of the neutron star and the line-of-sight \mathbf{n} . This is conceptually the simplest source model, and it does not entail any loss of generality, as the mapping between $\{h_0, \cos \iota\}$ and $A_{+, \times}$ is one-to-one.

4.2 Signals in noise

In practice, the strain $x(t)$ measured by a detector is mainly dominated by noise $n(t)$, such that even in the presence of a signal $h(t)$ we have

$$x(t) = n(t) + h(t; \mathcal{A}, \boldsymbol{\lambda}). \quad (31)$$

The measured output is not a continuous function of time, but a discrete time series of data points $x_i = x(t_i)$ with $t_i = i \Delta t$, sampled at a finite rate $f_{\text{samp}} = 1/\Delta t$, e.g. for LIGO and GEO600 $f_{\text{samp}} = 16,384$ Hz and $f_{\text{samp}} = 20,000$ Hz for Virgo. A continuous-time formulation is often used for convenience of notation, however. Let us make the idealized assumption of Gaussian stationary noise n_i with zero mean and covariance $\gamma_{jk} \equiv E[n_j n_k]$, where $E[\dots]$ denotes the expectation-value of a random variable. If we define the scalar product $(x||y)$ of two (real-valued) time series x_i and y_i as

$$(x||y) \equiv \sum_{j,l} x_j \gamma_{jl}^{-1} y_l, \quad (32)$$

then the probability of a particular (pure-noise) time series $\{n_j\}_{j=0}^{M-1}$ is expressible as

$$P(\{n_j\}|\gamma) = (2\pi)^{-M/2} |\gamma|^{-1/2} e^{-\frac{1}{2}(n||n)}. \quad (33)$$

In the continuum limit of $\Delta t \rightarrow 0$, the scalar product (32) can be shown (cf. [41]) to converge to the expression

$$(x||y) \xrightarrow{\Delta t \rightarrow 0} 4\Re \int_0^\infty \frac{\tilde{x}(f) \tilde{y}^*(f)}{S_n(f)} df, \quad (34)$$

which is the classical Wiener filter of matched-filtering theory [86]. Here $\tilde{x}(f)$ is the Fourier transform of $x(t)$, and $*$ denotes complex conjugation. $S_n(f)$ is the single-sided power spectral density, which is defined as the Fourier transform of the auto-correlation function, i.e.

$$S_n(f) = 2 \int_{-\infty}^{\infty} E[n(0) n(t)] e^{-i2\pi f t} dt. \quad (35)$$

In practice, this definition is not very useful for computing S_n , however, and a more practical estimate for S_n can be obtained from the finite discrete time series $\{n_j\}$ using the Wiener-Khinchine theorem, namely

$$S_n(f) \approx \frac{2}{T_{\text{obs}}} E [|\tilde{n}(f)|^2], \quad (36)$$

which holds exactly in the limit of infinite duration T_{obs} of the time-series, i.e. when $T_{\text{obs}} \equiv M \Delta t \rightarrow \infty$. Here $\tilde{n}(f)$ is the discrete Fourier transform of n_j , defined as

$$\tilde{n}(f) = \Delta t \sum_{j=0}^{M-1} n_j e^{-i2\pi f j \Delta t}. \quad (37)$$

The quantity $\sqrt{S_n(f)}$ (which has units of $\text{Hz}^{-1/2}$) is the most commonly-used measure of the noise performance of GW detectors, e.g. see Fig. 2. Note that in the case of nearly-monochromatic signals $h(t)$, as considered here, only a very narrow frequency band around the signal frequency f_0 will contribute to the scalar product (34). In this case, $S_n(f)$ can be approximated as constant in the neighborhood of f_0 , and the scalar product therefore simplifies to

$$(x||y) \approx \frac{2}{S_n(f_0)} \int_0^{T_{\text{obs}}} x(t) y(t) dt. \quad (38)$$

Using (31) and (33), the probability of measuring a strain $x(t)$ in the presence of Gaussian noise $n(t)$ and a signal $h(t; \mathcal{A}, \boldsymbol{\lambda})$ can be expressed as

$$P(x|\mathcal{A}, \boldsymbol{\lambda}, S_n) \propto e^{-\frac{1}{2}(x-h||x-h)}. \quad (39)$$

The signal-to-noise ratio (SNR) of matched filtering is conventionally defined as $\rho \equiv \sqrt{(h||h)}$, and for continuous GWs it is straightforward to show [e.g. 49, 38, 65] that the SNR scales as

$$\rho = \sqrt{(h||h)} \propto \frac{h_0}{\sqrt{S_n}} \sqrt{T_{\text{obs}} \mathcal{N}}, \quad (40)$$

with the observation time T_{obs} and the number of equal-sensitivity detectors \mathcal{N} . This illustrates why it is essential for continuous-wave searches to integrate the data over the longest possible observation time T_{obs} (and use as many sensitive detectors \mathcal{N} as possible).

There are two different ways of proceeding from this point, depending on the paradigm of statistics used: Bayesian or frequentist. These approaches yield sometimes similar-looking answers (especially in Gaussian noise), but they require fundamentally different interpretations and provide different tools in practice. The conceptual difference between the two frameworks lies in the meaning of “probability”, while the axioms for *calculating* with probabilities are the same in both cases.

4.3 Frequentist framework: hypothesis testing

The frequentist approach is based on the frequency-interpretation of probability: the probability $P(A)$ of an event A is defined as the limiting fraction of events A in an infinite number of “identical”¹ trials. The frequentist detection problem is formulated as one of hypothesis testing: let H_0 be the hypothesis that there is no signal in the data (i.e. $h = 0$), and H_1 stands for the hypothesis that there is a nonzero signal h . Now we need a criterion to decide in an optimal way between the two hypotheses given a measurement $x(t)$. This can be achieved by computing a scalar *detection statistic* $\Lambda(x)$, and setting a threshold Λ^* such that H_0 is accepted if $\Lambda(x) < \Lambda^*$, while H_1 is accepted otherwise. From the probability distribution (39) of $x(t)$, we can calculate $P(\Lambda|H_0)$ and $P(\Lambda|H_1)$ for the two hypothesis. This allows us to define the *false alarm* probability $f_A(\Lambda^*)$, namely

$$f_A(\Lambda^*) \equiv \int_{\Lambda^*}^{\infty} P(\Lambda|H_0) d\Lambda, \quad (41)$$

which is the probability of Λ crossing the threshold Λ^* despite H_0 being true. Similarly, we define the *false dismissal* probability $f_D(\Lambda^*, h)$ of a signal h as

$$f_D(\Lambda^*, h) \equiv \int_{-\infty}^{\Lambda^*} P(\Lambda|H_1) d\Lambda, \quad (42)$$

which is the probability that Λ does not cross the threshold Λ^* , even though H_1 is true. The detection probability $\eta \equiv \int_{\Lambda^*}^{\infty} \mathcal{P}(\Lambda|H_1) d\Lambda$ is then simply the complement to f_D , namely $\eta = 1 - f_D$. A standard criterion for the optimality of a hypothesis test $\Lambda(x)$ is that the test should maximize the detection probability $\eta(\Lambda^*, h)$ at a given false-alarm rate $f_A(\Lambda^*)$. According to the Neyman-Pearson lemma, the optimal test is given by the likelihood ratio, which is defined as

$$\Lambda(x; h) \equiv \frac{P(x|H_1)}{P(x|H_0)}. \quad (43)$$

Applying this to the Gaussian detection problem (39), we find

¹ The trials can obviously not be truly identical or they would yield the same result.

$$\ln \Lambda(x; h) = (x \| h) - \frac{1}{2}(h \| h), \quad (44)$$

which is the well-known expression for the matched-filtering amplitude. If some of the parameters of the signal $h(t; \mathcal{A}, \boldsymbol{\lambda})$ are unknown, one tries to find the *maximum* of $\ln \Lambda$ as a function of $\{\mathcal{A}, \boldsymbol{\lambda}\}$, which yields the corresponding maximum-likelihood estimators (MLE) for these parameters.

The \mathcal{F} -statistic

In targeted searches of GWs from known pulsars, the Doppler parameters $\boldsymbol{\lambda}$, i.e. sky-position \boldsymbol{n} and spins $f^{(k)}$, are usually well known, but even in this case one has generally no information about the four amplitude parameters \mathcal{A} , for which we have to find the maximum-likelihood estimators. In wide-parameter searches the Doppler-parameters $\boldsymbol{\lambda}$ of possible signals are also treated as unknown. However, as shown in [49], the maximization over \mathcal{A}^μ can be performed analytically, thereby reducing the number of remaining unknown parameters that need to be searched. Inserting the factored expression (27) for the filters $h(t; \mathcal{A}, \boldsymbol{\lambda})$ into (44), we find

$$\ln \Lambda(x; \mathcal{A}, \boldsymbol{\lambda}) = \mathcal{A}^\mu x_\mu - \frac{1}{2} \mathcal{A}^\mu \mathcal{A}^\nu \mathcal{M}_{\mu\nu}, \quad (45)$$

with implicit summation over $\mu, \nu = 1, 2, 3, 4$, and where we defined

$$x_\mu(\boldsymbol{\lambda}) \equiv (x \| h_\mu), \quad \text{and} \quad \mathcal{M}_{\mu\nu}(\boldsymbol{\lambda}) \equiv (h_\mu \| h_\nu). \quad (46)$$

We can now maximize $\ln \Lambda$ over \mathcal{A}^μ in order to obtain their maximum-likelihood estimators $\mathcal{A}_{\text{ML}}^\mu$ from the data $x(t)$, namely

$$\frac{\partial \ln \Lambda}{\partial \mathcal{A}^\mu} = 0 \quad \implies \quad \mathcal{A}_{\text{ML}}^\mu = \mathcal{M}^{\mu\nu} x_\nu, \quad (47)$$

where $\mathcal{M}^{\mu\alpha} \mathcal{M}_{\alpha\nu} = \delta_\nu^\mu$. Substituting these $\mathcal{A}_{\text{ML}}^\mu$ into (45), we obtain a new detection statistic, which only depends on the Doppler parameters $\boldsymbol{\lambda}$, namely

$$2\mathcal{F}(x; \boldsymbol{\lambda}) = x_\mu \mathcal{M}^{\mu\nu} x_\nu, \quad (48)$$

which is known as the “ \mathcal{F} -statistic” in this context². The coherent multi-detector generalization of the \mathcal{F} -statistic was derived recently by [38], but for simplicity we restrict the following discussion to the case of a single detector. In the presence of a signal $h(t; \mathcal{A}_s, \boldsymbol{\lambda}_s)$, the expectation value of the \mathcal{F} -statistic with perfectly matched Doppler parameters, $\boldsymbol{\lambda} = \boldsymbol{\lambda}_s$, is found as

$$E[2\mathcal{F}] = 4 + \rho^2, \quad (49)$$

where ρ is the optimal SNR defined previously in (40). One can show (cf. [49, 38]) that $2\mathcal{F}$ is a random variable with a χ^2 -distribution with 4 degrees

² Not to be confused with the F-statistic or the F-test in the statistics literature

of freedom and a non-centrality parameter ρ^2 (for the definition of the non-central χ^2 -distribution, see for example [10]). In the absence of a signal, i.e. $\rho = 0$, this reduces to the central χ^2 -distribution, namely

$$P(2\mathcal{F}; 0) = \frac{1}{2} \mathcal{F} e^{-\mathcal{F}}. \quad (50)$$

Using the known probability distribution of $2\mathcal{F}$, we can compute the false-alarm probability $f_A(2\mathcal{F}^*)$ and the false-dismissal probability $f_D(2\mathcal{F}^*; \rho^2)$ for a threshold $2\mathcal{F}^*$ and SNR ρ . The false-alarm probability (41) is easily integrated and yields

$$f_A(2\mathcal{F}^*) = (1 + \mathcal{F}^*) e^{-\mathcal{F}^*}, \quad (51)$$

while the false-dismissal probability (42) generally needs to be integrated numerically. If we choose a false-alarm rate of $f_A(2\mathcal{F}^*) = 1\%$, say, then (51) determines a detection threshold of $2\mathcal{F}^* \approx 13.3$. With this threshold, the required SNR for a false-dismissal rate of 10%, say, is given by the solution of $f_D(2\mathcal{F}^*; \rho^2) = 10\%$, which results in $\rho \approx 4.5$. Using Eq. (93) of [49] relating the *average* SNR $\langle \rho \rangle$ (over sky-location \mathbf{n} , orientation $\cos \iota$ and polarization ψ), to the amplitude h_0 , we can express the smallest average amplitude $\langle h_0 \rangle$ that would be detectable with a 1% false-alarm and 10% false-dismissal rate as

$$\langle h_0 \rangle_{f_D=10\%}^{f_A=1\%} \approx 11.4 \sqrt{\frac{S_n}{T_{\text{obs}}}}, \quad (52)$$

where T_{obs} is the coherently-integrated observation time. This is a useful measure of the sensitivity of a search. Note, however, that the false-alarm rate (51) refers to a *single* trial, and therefore this sensitivity-estimate only applies to targeted, single-template searches. When using a coherent network of detectors with respective noise floors S_X , the combined noise-floor to use in (52) is given by $S_n^{-1} = \sum S_X^{-1}$. Therefore, the combined sensitivity of \mathcal{N} equal-noise detectors is improved by a factor of $\sqrt{\mathcal{N}}$. In the case of LIGO I, the combined design-sensitivity of H1, L1 and H2 (cf. Sec. 3) would be roughly a factor $\sqrt{2.25}$ better, as H2 has half the arm length and is only about half as sensitive as H1 and L1. Advanced LIGO will consist of three 4 km interferometers, and the combined H1+H2+L1 detector will therefore be a factor $\sqrt{3}$ more sensitive than any single one.

Another important quantity is the *upper limit* on the amplitude of gravitational waves that we can obtain from an observation that did not detect a signal. The standard frequentist upper limit of confidence C is defined as the amplitude h_0^C of signals that would result in values $2\mathcal{F}$ exceeding the loudest candidate $2\mathcal{F}_0$ that was actually measured in a fraction C of trials, i.e.

$$C = \int_{2\mathcal{F}_0}^{\infty} P(2\mathcal{F}|h_0^C) d2\mathcal{F}, \quad (53)$$

which will often be computed using Monte-Carlo integration. Note that, contrary to the Bayesian approach described below, this is not a statement about

the confidence that the true value of h_0 is contained in the interval $[0, h_0^C]$, but about the frequency with which the so-constructed interval would contain the true value in repeated experiments. For a more detailed discussion and an elegant method of constructing frequentist confidence intervals, see [40].

4.4 Bayesian analysis: parameter estimation

Bayesian statistics is built on a different concept of probability, quantifying the degree of certainty (or “degree of belief”) of a statement being true (see [53, 71] for general introduction and references). Freed from the narrow frequentist definition of probability, one can now assign probabilities $P(A|I) \in [0, 1]$ to any statement A within some model I , quantifying one’s (possibly incomplete) knowledge about the truth of A . The probabilities $P(A|I) = 1$ and $P(A|I) = 0$ reflect the extreme cases of certainty about A being true or false, respectively. The axioms of probability provide a natural framework to calculate with such quantified partial knowledge. A central tool in this approach is Bayes’ theorem (an elementary consequence of the axioms of probability), namely

$$P(A|x, I) = P(x|A, I) \frac{P(A|I)}{P(x|I)}. \quad (54)$$

Using this together with (39), we can express the “posterior probability” of a signal $h(t; \mathcal{A}, \boldsymbol{\lambda})$ being present, *given* the measurement $x(t)$, as

$$P(\mathcal{A}, \boldsymbol{\lambda}|x, I) = k P(x|\mathcal{A}, \boldsymbol{\lambda}, I) P(\mathcal{A}, \boldsymbol{\lambda}|I), \quad (55)$$

where k is a normalization constant. The term $P(\mathcal{A}, \boldsymbol{\lambda}|I)$ is the “prior probability”, which expresses our previous knowledge about the signal, either from other measurements or from theoretical considerations. Eq. (55) quantifies how our state of knowledge is transformed from the prior $P(\mathcal{A}, \boldsymbol{\lambda}|I)$ to the posterior $P(\mathcal{A}, \boldsymbol{\lambda}|x, I)$ in the light of new information $x(t)$. One of the somewhat controversial aspects of Bayesian statistics is the assignment of prior probabilities. One often tries to use a prior reflecting “ignorance” or minimal bias (a common choice is a flat prior, i.e. $P(\mathcal{A}, \boldsymbol{\lambda}|I) = \text{const.}$), but this is not unproblematic and there is no unique choice of such a zero-information prior. Substituting the likelihood (39) in (55), we find the posterior probability

$$P(\mathcal{A}, \boldsymbol{\lambda}|x, I) = k' P(\mathcal{A}, \boldsymbol{\lambda}|I) \Lambda(x; \mathcal{A}, \boldsymbol{\lambda}), \quad (56)$$

where k' is another normalization constant. We see that are naturally lead back to an expression resembling the likelihood ratio (44). However, the statistical interpretation in this case is very different from the frequentist framework, as the posterior (56) determines the probability of a signal with certain parameters being present, while we cannot assign meaningful frequentist probabilities to such statements. Equation (56) formulates the detection problem in terms of *parameter estimation* of the signal, while the frequentist approach is typically based on *hypothesis testing*.

A powerful tool of the Bayesian framework is *marginalization* over “nuisance parameters”. Let us assume for simplicity that we have performed a targeted search with known Doppler parameters λ . If we want to express the posterior for only a subset of the four unknown amplitude parameters \mathcal{A} , e.g. we might be most interested in $P(h_0|x, I)$, say, then we can compute this by simply summing the posterior over all the possible values of the “uninteresting” parameters $\cos \iota$, ψ and ϕ_0 , i.e.

$$P(h_0|x, I) \propto \int P(\mathcal{A}|x, I) d\phi_0 d\psi d\cos \iota. \quad (57)$$

The Bayesian way of determining an upper limit $h_0^{\mathbf{C}}$ of confidence \mathbf{C} consists simply of solving the equation

$$\mathbf{C} = \int_0^{h_0^{\mathbf{C}}} P(h_0|x, I) dh_0, \quad (58)$$

i.e. the true amplitude h_0 lies within $[0, h_0^{\mathbf{C}}]$ with probability \mathbf{C} . Note that this “confidence” \mathbf{C} has an entirely different meaning from the frequentist confidence in Eq. (53). A more detailed discussion of the application of Bayesian statistics to the problem of detecting continuous GWs is found in [39] and [32].

4.5 Parameter space of coherent wide-parameter searches

The sensitivity of wide-parameter (as opposed to targeted) searches is severely limited by the large number of required templates, which is a rapidly growing function of the coherent observation time T_{obs} . The measure for how densely templates need to be placed in the Doppler-parameter space is determined by the relative loss m of detection statistic \mathcal{F} caused by an offset $d\lambda = \lambda - \lambda_s$ from a putative signal position λ_s . This mismatch m induces a natural distance measure and a corresponding local *metric* g_{ij} on the parameter space, first introduced by [21] and [59], namely

$$m(\lambda_s, d\lambda) = \frac{E[\mathcal{F}(\lambda_s)] - E[\mathcal{F}(\lambda)]}{E[\mathcal{F}(\lambda_s)]} = g_{ij}(\lambda_s) d\lambda^i d\lambda^j + \mathcal{O}(d\lambda^3), \quad (59)$$

where we used the fact that $E[\mathcal{F}(\lambda_s)]$ is a local maximum of \mathcal{F} if there is a signal in λ_s . It can be shown (cf. [26, 65]) that this metric can be approximated as

$$g_{ij} \sim \langle \partial_i \phi \partial_j \phi \rangle - \langle \partial_i \phi \rangle \langle \partial_j \phi \rangle, \quad (60)$$

in terms of the signal phase $\phi(t)$, and where we defined $\partial_i = \partial/\partial\lambda^i$ and where $\langle \dots \rangle$ denotes the time-average over T_{obs} . Considering the explicit phase model (22), (24) for isolated neutron stars with one spindown, i.e. $\lambda^i \in \{\mathbf{n}, f, \dot{f}\}$, one can easily show [e.g. 65] the following dominant scaling relations

$$g_{\theta\theta} \propto f^2 T_{\text{obs}}^2 (V/c)^2, \quad g_{ff} \propto T_{\text{obs}}^2, \quad g_{jj} \propto T_{\text{obs}}^4, \quad (61)$$

where θ is the angular separation on the sky and $V/c \sim 10^{-4}$ is the maximal Doppler shift due to the orbital velocity V . The required number of templates dN_p per small parameter-space region $d^4\boldsymbol{\lambda} = d\Omega \times df \times d\dot{f}$ (with sky solid-angle $d\Omega$) is therefore

$$dN_p \propto \sqrt{|\det g_{ij}|} d^4\boldsymbol{\lambda} \propto T_{\text{obs}}^5 f^2 d^4\boldsymbol{\lambda}. \quad (62)$$

As noted in [65], the $\mathcal{O}(T_{\text{obs}}^2)$ growth of the number of sky templates should only be considered as a lower bound, and a more detailed analysis is required to determine the exact scaling with T_{obs} . In a wide-parameter search we need to integrate (38) for each template, i.e. the computing cost c_p per template is $c_p \propto T_{\text{obs}}$, and the total computing cost dC_p per parameter space $d^4\boldsymbol{\lambda}$ scales (at least) as

$$dC_p \propto T_{\text{obs}}^6 f^2 d^4\boldsymbol{\lambda}. \quad (63)$$

On the other hand, recall that longer integration times are required in order to increase the SNR (40) of weak signals. This rapid growth of the computing cost $C_p(T_{\text{obs}})$ severely limits the longest integration time T_{obs} that is affordable for all-sky, wide-frequency searches to about $T_{\text{obs}} \sim 10 \text{ h} - 30 \text{ h}$ with currently realistic levels of computing power, e.g. ~ 50 Tflops with Einstein@Home (cf. Sec. 4.7). The second effect of the large number of templates N_p is to reduce the sensitivity compared to a targeted search with the *same* observation time and false-alarm probability: increasing the number of templates increases the number of expected false-alarm candidates at fixed detection threshold. Therefore the detection-threshold needs to be raised to maintain the same false-alarm rate, thereby decreasing the sensitivity.

Note that increasing the number of equal-sensitivity detectors \mathcal{N} improves the SNR (40) in the same way as increasing the integration time T_{obs} . However, the expression (60) for the metric and the scaling (61) are still approximately valid even for such a network of detectors, as shown in [65]. This implies that increasing the number of detectors \mathcal{N} does *not* increase the required number of templates N_p , which makes this the computationally cheapest way to improve the SNR of coherent wide-parameter searches.

4.6 Semi-coherent methods

Coherent matched-filtering is the optimal method for targeted, single-template searches. However, the discussion in the previous section shows that wide-parameter searches will require techniques that trade off statistical “optimality” for lower computing cost (i.e. a smaller number of templates). These “semi-coherent” methods are less sensitive than matched filtering for the same observation time. However, as they require far fewer templates in parameter space, they allow one to over-compensate this apparent loss of sensitivity by

using a *longer* observation time and *lower* thresholds, which typically results in more sensitive searches at substantially lower cost in computing power.

Here we focus on the simplest type of semi-coherent methods, which operate on successive short Fourier transforms (SFTs) of the measured strain data $x(t)$. There are three main variants of such SFT-based semi-coherent methods currently in use for GW searches, known respectively as “StackSlide” (also known as the “Radon transform”), the “Hough transform” and “PowerFlux” (see [6] for a more detailed description). The aim of these methods is to detect a statistical excess of power in the frequency-bins corresponding to the time-frequency “path” of the signal of frequency $\hat{f}(t)$ at the detector, as illustrated in Fig. 3. The instantaneous frequency $\hat{f}(t)$ of a signal at the detector can be




Fig. 3. An illustration of the principle used in semi-coherent methods to detect continuous GWs: The dark pixels represent a signal in the data. Its frequency changes with time due to Doppler shifts and intrinsic evolution of the source. By appropriately sliding the frequency bins of successive coherent “stacks” (shown vertically), the power of the signal can be lined up and summed (from [6]).

approximately (e.g. see [52]) related to the instantaneous intrinsic frequency $f(t)$ of a GW by

$$\hat{f}(t) = \left(1 + \frac{\mathbf{v}(t) \cdot \mathbf{n}}{c} \right) f(t), \quad (64)$$

where $\mathbf{v}(t)$ is the detector velocity with respect to the SSB frame. In the semi-coherent methods, the total observation time T_{obs} is divided into N “stacks” of duration $T_{\text{coh}} = T_{\text{obs}}/N$. The timeseries $x_j^{(l)}$ in stack l is Fourier-transformed, which gives the SFT $\tilde{x}_k^{(l)}$ for the stack l , with frequency bins $f_k = k/T_{\text{coh}}$. The normalized power $\rho_k^{(l)}$ in frequency-bin k of stack l is defined as

$$\rho_k^{(l)} \equiv \frac{2|\tilde{x}_k^{(l)}|^2}{T_{\text{coh}} S_n}, \quad (65)$$

such that in the absence of a signal, the expectation value is $E[\rho_k^{(l)}] = 1$. The maximal length of SFT stacks is constrained by the requirement that the signal power should not be spread over more than one frequency bin by the Doppler shift. This typically limits SFT stacks to about $T_{\text{coh}} \lesssim 60$ min.

The *StackSlide* method computes the total power $P = N^{-1} \sum_l \rho_{k(l)}^{(l)}$ along the path of frequency bins $k(l)$ corresponding to the signal frequency (64), as illustrated in Fig. 3. The *PowerFlux* method is a variant of StackSlide, summing *weighted* power in order to improve the sensitivity by taking account of non-stationarities of the noise and the direction-dependent antenna-patterns. The *Hough transform*, on the other hand, sums binary *number counts* $n_k^{(l)}$ instead of power. These number counts are obtained by setting a threshold ρ_{th} on the normalized power (65), namely $n_k^{(l)} = 1$ if $\rho_k^{(l)} \geq \rho_{\text{th}}$, and $n_k^{(l)} = 0$ otherwise. The final Hough detection statistic is the total number count $n = \sum_l n_{k(l)}^{(l)}$ of threshold-crossings along the time-frequency path $k(l)$. Summing binary number counts instead of power slightly reduces the sensitivity of the Hough method compared to StackSlide and PowerFlux, but it increases its robustness with respect to transient disturbances, as no single stack can contribute more than ‘1’ to the final number count.

Semi-coherent methods combine “power” from different stacks, while the phase information between stacks has been lost, which reduces the sensitivity compared to fully coherent matched filtering with the same observation time T_{obs} (by increasing the false-alarm probability). This effect can be seen, for example, in the expression for the average amplitude h_0 of the weakest signal detectable in a targeted Hough-search with a false-alarm rate of $f_A = 1\%$ and a false-dismissal of $f_D = 10\%$ [52], namely

$$\langle h_0 \rangle_{f_A=1\%}^{f_D=10\%} \approx 8.5 N^{1/4} \sqrt{\frac{S_n}{T_{\text{obs}}}} = 8.5 N^{-1/4} \sqrt{\frac{S_n}{T_{\text{coh}}}}. \quad (66)$$

Similar expressions hold for StackSlide and PowerFlux [6]. Comparing this to the analogous matched-filtering expression (52) shows that a coherent search over the full observation time T_{obs} would be more sensitive by about a factor of $N^{1/4}$. This comparative loss in sensitivity, however, is accompanied by an enormous advantage, namely a substantially lower parameter-space resolution compared to a coherent search with the same T_{obs} . We can estimate the frequency resolution of a semi-coherent search as $\delta f \sim 1/T_{\text{coh}}$, which is the frequency-resolution of the SFT stacks. The resolution in spindown $\delta \dot{f}$ and angular sky-position $\delta \theta$ can be estimated from the requirement that the frequency should not drift by more than one frequency-bin δf over the total observation time T_{obs} (cf. [6]), which results in

$$\delta \dot{f} \sim \frac{1}{T_{\text{obs}} T_{\text{coh}}}, \quad \delta \theta \sim \frac{1}{f T_{\text{coh}} V/c}. \quad (67)$$

The number of required templates per parameter-space $d^4 \boldsymbol{\lambda}$ in a search for isolated neutron stars with one spindown therefore scales as

$$dN_p \propto T_{\text{obs}} T_{\text{coh}}^4 f^2 d^4 \boldsymbol{\lambda}. \quad (68)$$

Each parameter-space point requires summing $N \propto T_{\text{obs}}$ numbers, so the computing cost dC_p scales as

$$dC_p \propto T_{\text{obs}}^2 T_{\text{coh}}^4 f^2 d^4 \lambda, \quad (69)$$

which shows that the increase in computing cost with T_{obs} is *substantially* weaker than for the fully coherent case (63). The search can therefore be extended over much longer total observation times T_{obs} , of the order of several months, thereby achieving a better sensitivity at lower computing cost than fully coherent matched-filtering. The semi-coherent methods are not restricted to using SFT stacks, but we can also use “demodulated” stacks [52] such as the \mathcal{F} -statistic (48). This allows to increase the length of the coherent stacks T_{coh} beyond the short duration $T_{\text{coh}} \lesssim 60$ min of SFT stacks, which increases the sensitivity (66), but also the computing cost (69).

4.7 Hierarchical searches and Einstein@Home

Neither the matched-filtering nor the semi-coherent methods described in the previous sections optimize by themselves the sensitivity of wide-parameter searches at given finite computing power: the sensitivity can be further improved by appropriately combining several stages of such coherent and semi-coherent steps, in what is generally known as “hierarchical” schemes [25, 42, 36].

The general idea is to start with a wide-parameter search using a relatively short observation time and therefore low resolution in parameter space. In this first stage a low threshold is used in order to increase the chances of a weak signal crossing the threshold. This will result in a large number of random-chance candidates, however, which are followed up in a second stage search using a longer observation time and a higher threshold. This is computationally affordable due to the reduced number of templates required to follow up the first-stage candidates, as opposed to scanning the entire parameter space at high resolution. This step can be iterated several times with increasingly longer observation times and higher thresholds, successively gaining confidence in the surviving candidates. There are several free parameters in such a scheme, such as the number of stages, the length of respective stacks and their corresponding thresholds, all of which need to be optimized in order to obtain the best possible sensitivity per computing cost. A first study of this optimization problem was carried out by [36], and the results suggest that about three stages might be sufficient.

In addition to the (still ongoing) effort to develop such an optimal search algorithm, one also wants to maximize the available computing power in order to optimize the absolute sensitivity of the search. This second goal is accomplished by the Einstein@Home project,³ a public distributed-computing project launched in Feb. 2005. Einstein@Home is based on the distributed-computing platform BOINC,⁴ which was originally developed for Seti@Home,

³ <http://einstein.phys.uwm.edu/>

⁴ <http://boinc.berkeley.edu/>

and which is now used by a growing number of distributed-computing projects. The search for continuous GWs is ideally suited for this kind of distributed approach, as it can be split into a large number of small, independent problems: each participating host analyzes only a small portion $\Delta\lambda$ of the total Doppler parameter space. After completing this search, the host returns the results to a central project server and requests the next “work-unit” to analyze. By Aug. 2006, Einstein@Home has attracted more than 100,000 participants, contributing more than 200,000 CPUs, and delivering more than 50 Tflops of continuous computing power. The hierarchical search scheme currently under development for Einstein@Home is ultimately expected to yield the most sensitive wide-parameter search available for continuous GWs from unknown spinning neutron stars.

5 Current status of the search for continuous GWs

5.1 Overview of continuous-wave searches

Figure 4 gives an overview of the current status of searches for continuous GWs from spinning neutron stars. This figure shows *approximate* levels of published upper limits, *estimated* sensitivities of current and future searches, and astrophysically motivated upper limits. The spindown upper limits (18) for known pulsars shown in Fig. 4 are based on the pulsar parameters in the ATNF catalog [54], allowing for a distance uncertainty of $\pm 10\%$ and a moments of inertia in the range $I_{zz} = (1 - 3) \times 10^{38} \text{ kg m}^2$ (see [8] for discussion). The astrophysical limit for Sco X-1 is based on the Bildsten-Wagoner mechanism (20). The curves ‘LIGO-I’, ‘Virgo’, ‘eLIGO’ and ‘AdvLIGO’ in Fig. 4 are based on (52), describing a targeted coherent integration over one year of data at the respective design sensitivity. The AdvLIGO curve corresponds to a wideband tuning of the detectors (with anticipated fundamental noise sources), but Advanced LIGO can also be tuned in different ways in order to pinpoint an anticipated high-frequency source with narrow-banding, or increase the sensitivity at low frequencies (at the expense of higher frequencies). In the case of LIGO-I, eLIGO and AdvLIGO the sensitivities plotted in Fig. 4 refer to multi-detector searches using the network of three detectors (H1+H2+L1). These sensitivities can only be reached in targeted searches for known pulsars, assuming ideal conditions of well-constrained parameters (e.g. negligible timing-noise and no glitches). As discussed in Sec. 4.5, wide-parameter searches are inherently less sensitive due to the large number of templates required.

5.2 Results from completed searches

S1: targeted search for PSR B1937+21

The fully coherent methods described in Sec. 4.3 and Sec. 4.4 were used to perform a targeted search for the millisecond pulsar PSR B1937+21, using

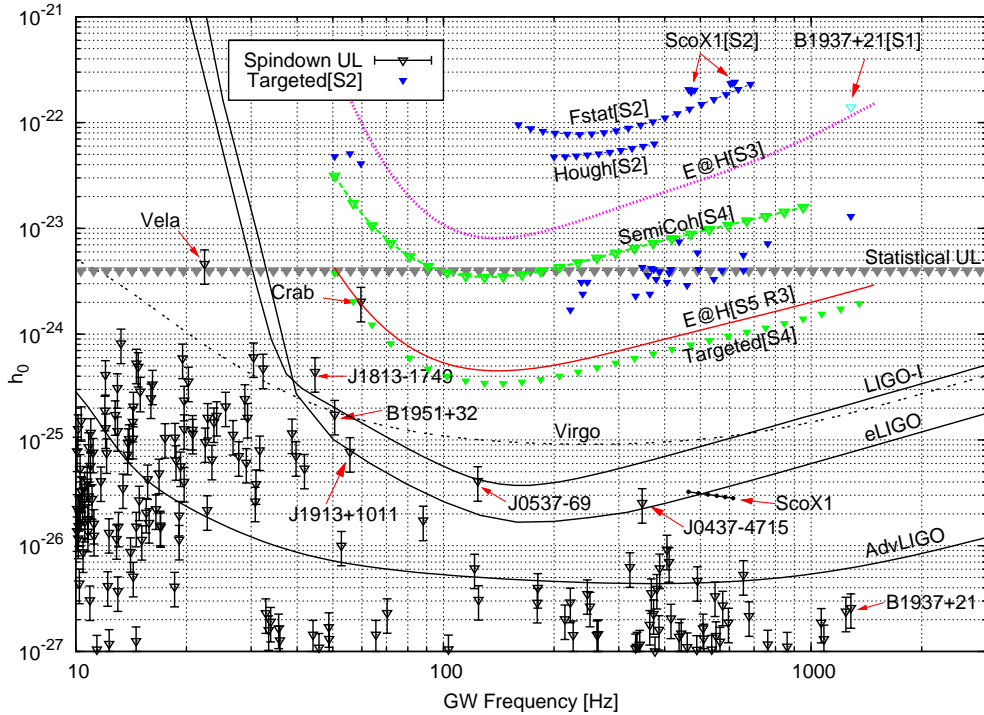


Fig. 4. Summary plot of approximate levels of published upper limits (cf. Sec. 5.2), estimated sensitivities of various current and future searches (cf. Sec. 5.3) and astrophysical upper limits (cf. Sec. 5.1). The ‘Statistical UL’ refers to the loudest expected signal from a population of unknown isolated neutron stars spinning down due to GWs (cf. Sec. 2.2, Eq. (14)).

data from the first LIGO science run (S1) [2]. The Doppler parameters (sky-position + spin) of this pulsar are well known and are extremely stable, which allows a single-template search. From the spindown upper limit for this pulsar (cf. Fig. 4), it is obvious that no detection was expected, the motivation for this search was mainly to illustrate the methods and set an upper limit on the GW emission at $f = 2\nu$. The resulting best upper limit obtained was $h_0^{95\%} \sim 1.4 \times 10^{-22}$.

S2: \mathcal{F} -statistic wide-parameter search for Sco X-1

A coherent wide-parameter \mathcal{F} -statistic search for Sco X-1 was performed [7] using data from the S2 science run. Sco X-1 is a neutron star in a 18.9 h orbit around a low-mass companion, at a distance of $d \sim 2.8$ kpc from Earth. The sky-position \mathbf{n} and orbital period P of Sco X-1 are well determined from X-ray observations, but both the projected semi-major axis a_p and the orbital

phase \bar{T} have large uncertainties and need to be treated as unknown (Doppler) parameters. The rotation frequency ν of the neutron star is also highly uncertain. Assuming the (not uncontroversial) beat-frequency model for QPOs [83], and a triaxial neutron star emitting GWs at a frequency $f = 2\nu$, the frequency window of the search would span at least $f \in [460, 620]$ Hz. Due to the enormous computational cost of this search, however, the frequency band had to be reduced to two smaller bands $f \in [464, 484]$ Hz and $f \in [604, 624]$ Hz. The scaling of the number of templates N_p for this Doppler parameter space ($\lambda = \{f, a_p, \bar{T}\}$) is $N_p \propto T_{\text{obs}}^6$, which severely limited the maximum possible observation time T_{obs} . The analysis pipeline consisted of two \mathcal{F} -statistic searches over the most sensitive $T_{\text{obs}} = 6$ h of data from L1 and H1 respectively, followed by a coincidence step to reduce the number of false-alarm candidates. Upper limits of the order $h_0^{95\%} \sim 2 \times 10^{-22}$ were obtained over the range of parameters analyzed (labeled ‘ScoX1[S2]’ in Fig. 4).

S2: \mathcal{F} -statistic all-sky search for unknown isolated neutron stars

A similar wide-parameter search with the \mathcal{F} -statistic was performed for unknown isolated neutron stars over the whole sky and in the frequency range $f \in [160, 729]$ Hz, assuming a frequency-derivative of less than $|\dot{f}| < 4 \times 10^{-10}$ Hz/s. The scaling of the number of templates with observation time T_{obs} for this Doppler space $\lambda = \{\alpha, \delta, f\}$ is $N_p \propto T_{\text{obs}}^3$, see Sec. 4.5. The analysis consisted of a similar pipeline to the Sco X-1 search, using the most sensitive $T_{\text{obs}} = 10$ h of data from L1 and H1, respectively, and reducing the false-alarm rate by a coincidence step. The best all-sky (frequentist) upper limit achieved was of the order $h_0^{95\%} \sim 7 \times 10^{-23}$, as shown in Fig. 4 (label ‘Fstat[S2]’).

S2: Hough all-sky search for unknown isolated neutron stars

The Hough-transform method (cf. Sec. 4.6) was used in an all-sky search for unknown isolated neutron stars [3] in the frequency range $f \in [200, 400]$ Hz, including one spindown parameter $\dot{f} \in [-10^{-9}, +10^{-10}]$ Hz/s. The data from all three LIGO detectors was analyzed over the *whole* duration of $T_{\text{obs}} = 2$ months of S2. This required searching of the order $N_p \sim 10^{12}$ templates (many orders of magnitude less than a coherent search would require for the same T_{obs}), which took less than half a day to complete on a 200-CPU cluster. The best all-sky upper limit obtained in this frequency-range was $h_0^{95\%} \sim 4.4 \times 10^{-23}$, and the results are summarized in Fig. 4 (label ‘Hough[S2]’). The sensitivity achieved by this semi-coherent search is about a factor of *two* better than the equivalent \mathcal{F} -statistic search (‘Fstat[S2]’), which required similar computing power. This is a consequence of the lower parameter-space resolution of semi-coherent methods, which allowed to use the full two months of data as opposed to only 10 hours (cf. Sec. 4.6).

S2: targeted search for 28 known pulsars

The results from a targeted, fully coherent search for 28 known isolated radio pulsars in the LIGO band ($f = 2\nu \gtrsim 50$ Hz) have been reported in [4]. The analysis was based on the Bayesian approach described in Sec. 4.4, but the actual implementation used a highly efficient complex-heterodyning method described in [39]. Data from the S2 science run was analyzed, combining H1, L1 and H2 coherently (GEO600 did not take part in S2). The resulting Bayesian 95% confidence upper limits are shown in Fig. 4 (label ‘Targeted[S2]’). Most of these upper limits are still a few orders of magnitude above the corresponding spindown limits (where available), but for the globular-cluster pulsars with apparent spin-up, these are the first direct constraints available. The best strain upper limit obtained was $h_0^{95\%} \sim 1.7 \times 10^{-24}$ (for PSR J1910-5959D), and the best upper limit on the quadrupolar deformation was $e^{95\%} \sim 4.5 \times 10^{-6}$ (for PSR J2124-3358). The upper limit for the Crab pulsar (PSR B0531+21) was found as $h_0^{95\%} \sim 4 \times 10^{-23}$, which is within a factor of 30 from the spindown upper limit. As discussed in Sec. 4.5, fully coherent targeted searches provide the best possible sensitivity, which is illustrated by these results: at similar computing cost, the targeted S2 search is about 25 times more sensitive than the semi-coherent searches, and about 50 times more sensitive than the fully coherent wide-parameter searches.

S4: Coherent targeted searches for known pulsars

Using the data from all three LIGO interferometers and GEO600 taken during the S3 and S4 science runs, a fully coherent targeted search for known pulsars (at $f = 2\nu$) has recently been completed [8]. The detailed upper limits are found in this paper, here we restrict ourselves to estimating the average sensitivity of this search by Eq. (52), and this estimate is shown in Fig. 4 (label ‘Targeted[S4]’). According to this rough estimate the search would already appear to beat the spindown limit of the Crab pulsar, but unfortunately this is not the case: the noise is less stationary at lower frequencies and there are also some side-effects from the strong spectral disturbance at 60 Hz (visible in Fig. 2), which stems from the mains power-line frequency. As a result the upper limit for the Crab pulsar still lies about a factor of 2.2 above the spindown limit in this search. Using several months of S5 data, however, the Crab spindown limit will be beaten for the first time, setting the astrophysically most relevant upper limit on the GW emission of any pulsar so far.

S4: Semi-coherent searches

An all-sky search for unknown isolated neutron stars in the frequency range $f \in [50, 1000]$ Hz on data from the S4 run ($T_{\text{obs}} \sim 500$ h) was completed [6] using the semi-coherent methods described in Sec. 4.6, namely Hough, StackSlide and PowerFlux. The stacks consisted of SFTs of duration $T_{\text{coh}} =$

30 min, and data from all three LIGO interferometers was used. The expected sensitivity for these searches can be directly estimated using Eq. (66), which is shown in Fig. 4 (label ‘SemiCoherent[S4]’).

5.3 Ongoing and future searches

Einstein@Home: S3 and beyond

Einstein@Home (cf. Sec. 4.7) has completed a search on $T_{\text{obs}} = 600$ h of LIGO S3 data, and the results have been posted online [5]. The analysis pipeline consisted of $N = 60$ stacks (of $T_{\text{coh}} = 10$ h each) of coherent all-sky, wide-frequency searches using the \mathcal{F} -statistic. This coherent step was performed on the participating hosts and the results were returned to the central server for post-processing, where they were combined using a stacking/coincidence scheme.

Einstein@Home is aiming for detection, and no upper limits have been set so far. In order to obtain a rough estimate of the *approximate* level of sensitivity of such searches, we assume (somewhat optimistically) that the post-processing step results in a $\propto N^{1/4}$ improvement in sensitivity, as is characteristic for semi-coherent methods, e.g. (66). However, the numerical prefactor in (66) critically depends on the threshold $2\mathcal{F}^*$ used in the coherent \mathcal{F} -statistic stacks. For Einstein@Home we had to choose a relatively high threshold of $2\mathcal{F}^* = 25$, in order to limit the total amount of (false-alarm) data from the individual stacks sent back to the central server for post-processing. Furthermore, due to the large total number of templates searched, we use a low false-alarm probability of $f_A = 3 \times 10^{-13}$. Using Eq. (6.40) of [52] and applying an extra factor of $1/\sqrt{m} \sim 1.4$ in order to account for an average grid mismatch of $m = 0.5$, we find a numerical prefactor of ~ 72 instead of 8.5 as used in (66). The corresponding sensitivity-estimate is plotted in Fig. 4 (label ‘E@H[S3]’). A similar search was performed on S4 data, using a total of $T_{\text{obs}} = 510$ h of data divided in $N = 17$ stacks of \mathcal{F} -statistic integrations over $T_{\text{coh}} = 30$ h. The results from this run are currently in the post-processing stage. A search on S5 data using the same pipeline has been completed as well, using $N = 28$ stacks of $T_{\text{coh}} = 30$ h.

The sensitivity of all these searches suffered from the same problem mentioned above for the S3 search, namely the high $2\mathcal{F}^*$ -threshold required. This limitation, however, will be overcome in the setup for the current Einstein@Home search (labeled ‘S5 R3’), which includes a semi-coherent Hough step *on* the host, combining the \mathcal{F} -statistic searches from 84 stacks of $T_{\text{coh}} = 25$ h. This hierarchical approach substantially reduces the amount of (false-alarm) data that needs to be sent back, allowing to use the optimal threshold $2\mathcal{F}^* = 5.2$ [52]. We can estimate the corresponding ‘optimal’ sensitivity of such an Einstein@Home search, which yields a numerical prefactor of ~ 20 in Eq. (66), and the corresponding sensitivity estimate is shown in Fig. 4, labeled ‘E@H [S5 R3]’.

5.4 Previous upper limits from other detectors

Bar detectors: an earlier attempt to specifically target the Crab pulsar (at $f = 2\nu \sim 60$ Hz) was made with a specially-designed bar detector [74], setting an upper limit of $h_0 \sim 2 \times 10^{-22}$. A search targeting the millisecond pulsar PSR B1937+21 (at $f = 2\nu \sim 1284$ Hz) was performed by [46] using a split bar detector, producing an upper limit of $h_0 \sim 10^{-20}$. A search for unknown isolated neutron stars in a small frequency-band $f = (921.35 \pm 0.03)$ Hz and a small sky-region in the galactic center was performed using 95 days of data from the EXPLORER bar detector, and an upper limit of $h_0 \sim 3 \times 10^{-24}$ was obtained [17]. An all-sky search with two days of EXPLORER data in the frequency-band $f = (921.38 \pm 0.38)$ Hz was carried out using the \mathcal{F} -statistic, setting an upper limit of $h_0^{99\%} \sim 2 \times 10^{-23}$, reported in [18]. This, was later revised to $h_0^{90\%} \sim 10^{-22}$ in [19] using different conventions to determine the upper-limit. Another directed search was presented in [55]: data from the ALLEGRO bar detector was searched for periodic GWs from the Galactic center and from the globular cluster 47 Tuc in the two antenna bands (896.80 ± 0.50) Hz and (920.26 ± 0.50) Hz, setting an upper limit of $h_0 \sim 8 \times 10^{-24}$.

Interferometers: the first search using a broadband interferometer was carried out with the prototype 40 m detector at Caltech by [45]. The search was targeting PSR B1937+21, and resulted in upper limits of the order $h_0 \sim 3 \times 10^{-17}$ at $f = \nu$, and $h_0 \sim 1.5 \times 10^{-17}$ at $f = 2\nu$. Data from the first science run of the TAMA detector was searched for continuous GWs from SN1987A using coherent matched filtering over $T_{\text{obs}} = 1200$ h of data in a 0.05 Hz band at ~ 934.9 Hz, and the reported upper limit [72] was about $h_0^{99\%} \sim 5 \times 10^{-23}$. An earlier upper-limit result on SN1987A was obtained from a directed search using $T_{\text{obs}} = 100$ h of data from the Garching prototype interferometer, which determined an upper limit of $h_0^{95\%} \sim 9 \times 10^{-21}$ in 4 Hz bands around $f = \nu \sim 1670$ Hz and $f = 2\nu$ [58].

6 Outlook

LIGO has made enormous progress over the past 4 years, as seen in Fig. 2, reaching its design sensitivity with the S5 science run. This progress in sensitivity is also reflected in improving sensitivities of successive continuous-wave searches shown in Fig. 4. A detection of a spinning neutron star with LIGO-I is still somewhat unlikely (albeit not implausible), but continuous-wave searches are already beginning to enter a regime of increasing astrophysical relevance:

(i) A targeted pulsar search with S5 will beat the spindown upper limit for the Crab pulsar. This is the first time that direct GW observations set an astrophysical constraint on the contribution of GWs to the observed pulsar spindown.

(ii) With one year of data from S5, the spindown limits of at least two more pulsars (J0537-6910 and B1951+32) should be (marginally) reachable (see Fig. 4).

(iii) The upper limits on the non-axisymmetry ϵ of known pulsars are entering the regime $\epsilon \lesssim 10^{-6}$, which is physically possible according to our current understanding of neutron-star physics (cf. Sec. 2.1).

(iv) The statistical upper-limit level $h_0 \sim 4 \times 10^{-24}$ for unknown isolated neutron stars (cf. Sec. 2.2) has been marginally reached by the semi-coherent searches using S4 data, and will be substantially surpassed by semi-coherent searches using S5 (especially Einstein@Home). Passing this milestone suggests that the possibility of a detection of an unknown isolated neutron star is becoming increasingly plausible.

Furthermore, the Virgo detector is beginning to approach comparable sensitivities to LIGO. Once it has reached its design sensitivity, Virgo should be able to beat the spindown limit of up to three more known pulsars (including Vela), as seen in Fig. 4. Combining Virgo with the LIGO detector network will result in a further increase of $\sim 20\%$ in average sensitivity with respect to LIGO-I.

The Enhanced LIGO (eLIGO) upgrade might allow to reach the spindown limit of two more known pulsars, and to beat the spindown limit of J0537-6910. Following this, Advanced LIGO will be by far the most exciting instrument in the near future for GW searches from neutron stars, as seen in Fig. 4. It will allow to surpass the spindown limits of several tens of known pulsars with one year of data, and it will comfortably include the Bildsten-Wagoner emission-level (20) of Sco X-1 (assuming, however, that substantially better observational constraints on the Sco X-1 parameters are available for a directed search). Advanced LIGO will dig down nearly three orders of magnitude below the statistic upper limit (14), making a detection of neutron-star signals seem rather likely. Given the current pace of progress and the encouraging prospects about future developments, it seems reasonable to be optimistic that gravitational-wave astrophysics of neutron stars will finally become a reality within the next decade or two.

Acknowledgement. The authors gratefully acknowledge the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory and the Particle Physics and Astronomy Research Council of the United Kingdom, the Max-Planck-Society and the State of Niedersachsen/Germany for support of the construction and operation of the GEO600 detector. The authors also gratefully acknowledge the support of the research by these agencies and by the Australian Research Council, the Natural Sciences and Engineering Research Council of Canada, the Council of Scientific and Industrial Research of India, the Department of Science and Technology of India, the Spanish Ministerio de Educacion y Ciencia, The National Aeronautics and Space Administration, the John Simon Guggenheim Foundation, the Alexander von Humboldt Foundation, the Leverhulme Trust, the David and Lucile Packard Foundation, the Research Corporation, and the Alfred P.

Sloan Foundation. This document has been assigned LIGO Laboratory document number LIGO-P060039-v3.

References

- [1] Abbott, B. et al. (LIGO Scientific Collaboration): 2004a, *Nucl. Instrum. Methods Phys. Res. A* **517**, 154
- [2] Abbott, B. et al. (LIGO Scientific Collaboration): 2004b, *Phys. Rev. D.* **69**, 082004
- [3] Abbott, B. et al. (LIGO Scientific Collaboration): 2005a, *Phys. Rev. D.* **72**, 102004
- [4] Abbott, B. et al. (LIGO Scientific Collaboration): 2005b, *Phys. Rev. Lett.* **94**, 181103
- [5] Abbott, B. et al. (LIGO Scientific Collaboration): 2005c, *Report on Einstein@Home S3 analysis*, <http://einstein.phys.uwm.edu/>
- [6] Abbott, B. et al. (LIGO Scientific Collaboration): 2008, *Phys. Rev. D.* **77**, 022001
- [7] Abbott, B. et al. (LIGO Scientific Collaboration): 2007, *Phys. Rev. D.* **76**, 082001
- [8] Abbott, B. et al. (LIGO Scientific Collaboration), Kramer, M., and Lyne, A. G.: 2007, *Phys. Rev. D.* **76**, 042001
- [9] Abramovici, A. et al.: 1992, *Science* **256**, 325
- [10] Abramowitz, M. and Stegun, I. A.: 1964, *Handbook of Mathematical Functions*, National Bureau of Standards
- [11] Andersson, N.: 1998, *ApJ* **502**, 708
- [12] Andersson, N.: 2003, *Class. Quant. Grav.* **20**, 105
- [13] Andersson, N., Jones, D. I., and Kokkotas, K. D.: 2002, *MNRAS* **337**, 1224
- [14] Andersson, N. and Kokkotas, K. D.: 1998, *MNRAS* **299**, 1059
- [15] Andersson, N. and Kokkotas, K. D.: 2001, *International Journal of Modern Physics D* **10**, 381
- [16] Andersson, N., Kokkotas, K. D., and Stergioulas, N.: 1999, *ApJ* **516**, 307
- [17] Astone, P. et al.: 2002, *Phys. Rev. D.* **65**, 022001
- [18] Astone, P. et al.: 2003, *Class. Quant. Grav.* **20**, 665
- [19] Astone, P. et al.: 2005, *Class. Quant. Grav.* **22**, S1243
- [20] Baiotti, L. et al.: 2007, *Phys. Rev. D.* **75**, 044023
- [21] Balasubramanian, R., Sathyaprakash, B. S., and Dhurandhar, S. V.: 1996, *Phys. Rev. D.* **53**, 3033
- [22] Bildsten, L.: 1998, *ApJ* **501**, L89
- [23] Bonazzola, S. and Gourgoulhon, E.: 1996, *A&A* **312**, 675
- [24] Bonazzola, S. and Gourgoulhon, E.: 1997, in J.-A. Marck and J.-P. Lasota (eds.), *Relativistic Gravitation and Gravitational Radiation*, p. 151
- [25] Brady, P. R. and Creighton, T.: 2000, *Phys. Rev. D.* **61**, 082001

- [26] Brady, P. R., Creighton, T., Cutler, C., and Schutz, B. F.: 1998, *Phys. Rev. D* **57**, 2101
- [27] Caron, B., Dominjon, A., Drezen, C., Flaminio, R., et al.: 1997, *Nuclear Physics B Proceedings Supplements* **54**, 167
- [28] Carter, B., Langlois, D., and Sedrakian, D. M.: 2000, *A&A* **361**, 795
- [29] Centrella, J. M., New, K. C. B., Lowe, L. L., and Brown, J. D.: 2001, *ApJL* **550**, L193
- [30] Chakrabarty, D. et al.: 2003, *Nature* **424**, 42
- [31] Chandrasekhar, S.: 1970, *Phys. Rev. Lett.* **24**, 611
- [32] Christensen, N., Dupuis, R. J., Woan, G., and Meyer, R.: 2004, *Phys. Rev. D* **70**, 022001
- [33] Clark, J., Heng, I. S., Pitkin, M., and Woan, G.: 2007, *Phys. Rev. D* **76**, 043003
- [34] Cook, G. B., Shapiro, S. L., and Teukolsky, S. A.: 1994, *ApJ* **424**, 823
- [35] Cutler, C.: 2002, *Phys. Rev. D* **66**, 084025
- [36] Cutler, C., Gholami, I., and Krishnan, B.: 2005, *Phys. Rev. D* **72**, 042004
- [37] Cutler, C. and Jones, D. I.: 2001, *Phys. Rev. D* **63**, 024002
- [38] Cutler, C. and Schutz, B. F.: 2005, *Phys. Rev. D* **72**, 063006
- [39] Dupuis, R. J. and Woan, G.: 2005, *Phys. Rev. D* **72**, 102002
- [40] Feldman, G. J. and Cousins, R. D.: 1998, *Phys. Rev. D* **57**, 3873
- [41] Finn, L. S.: 1992, *Phys. Rev. D* **46**, 5236
- [42] Frasca, S., Astone, P., and Palomba, C.: 2005, *Class. Quant. Grav.* **22**, 1013
- [43] Friedman, J. L. and Morsink, S. M.: 1998, *ApJ* **502**, 714
- [44] Friedman, J. L. and Schutz, B. F.: 1978, *ApJ* **222**, 281
- [45] Hereld, M.: 1984, *Ph.D. Thesis (Caltech)*
- [46] Hough, J. et al.: 1983, *Nature* **303**, 216
- [47] Jaranowski, P. and Królak, A.: 1994, *Phys. Rev. D* **49**, 1723
- [48] Jaranowski, P. and Królak, A.: 2005, *Living Reviews in Relativity* **8**, 3, <http://www.livingreviews.org/lrr-2005-3>
- [49] Jaranowski, P., Królak, A., and Schutz, B. F.: 1998, *Phys. Rev. D* **58**, 063001
- [50] Jones, D. I. and Andersson, N.: 2002, *MNRAS* **331**, 203
- [51] Kittel, C.: 2005, *Introduction to solid state physics*, Wiley, 8th edition
- [52] Krishnan, B., Sintes, A. M., Papa, M. A., Schutz, B. F., Frasca, S., and Palomba, C.: 2004, *Phys. Rev. D* **70**, 082001
- [53] Loredó, J. T.: 1990, in P. F. Fougère (ed.), *Maximum Entropy and Bayesian Methods*, pp 81–142, Kluwer Academic Publishers, Netherlands
- [54] Manchester, R. N. et al.: 2005, *Astron. J.* **129**, 1993, <http://www.atnf.csiro.au/research/pulsar/psrcat/>
- [55] Mauceli, E. et al.: 2000, (preprint gr-qc/0007023)
- [56] Melatos, A. and Payne, D. J. B.: 2005, *ApJ* **623**, 1044
- [57] Misner, C. W., Thorne, K. S., and Wheeler, J. A.: 1973, *Gravitation*, W. H. Freeman and Company (New York)
- [58] Niebauer, T. M. et al.: 1993, *Phys. Rev. D* **47**, 3106

- [59] Owen, B. J.: 1996, *Phys. Rev. D.* **53**, 6749
- [60] Owen, B. J.: 2005, *Phys. Rev. Lett.* **95**, 211101
- [61] Owen, B. J.: 2006, *Class. Quant. Grav.* **23**, 1
- [62] Owen, B. J. et al.: 1998, *Phys. Rev. D.* **58**, 084020
- [63] Papaloizou, J. and Pringle, J. E.: 1978, *MNRAS* **184**, 501
- [64] Payne, D. J. B. and Melatos, A.: 2006, *ApJ* **641**, 471
- [65] Prix, R.: 2007, *Phys. Rev. D.* **75**, 023004
- [66] Ruderman, M.: 1969, *Nature* **223**, 597
- [67] Ruderman, M.: 1976, *ApJ* **203**, 213
- [68] Saijo, M. and Gourgoulhon, E.: 2006, *Phys. Rev. D.* **74**, 084006
- [69] Schutz, B. F. and Tinto, M.: 1987, *MNRAS* **224**, 131
- [70] Shapiro, S. L. and Zane, S.: 1998, *ApJS* **117**, 531
- [71] Sivia, D. S.: 1996, *Data Analysis. A Bayesian Tutorial*, Oxford Science Publications
- [72] Soida, K. et al. (TAMA Collaboration): 2003, *Class. Quant. Grav.* **20**, 645,
- [73] Stergioulas, N.: 2003, *Living Reviews in Relativity* **6**, 3, <http://relativity.livingreviews.org/Articles/lrr-2003-3/>
- [74] Suzuki, T.: 1995, in E. Coccia, G. Pizzella, and F. Ronga (eds.), *First Edoardo Amaldi Conference on Gravitational Wave Experiments*, p. 115
- [75] Taylor, J. H. and Weisberg, J. M.: 1989, *ApJ* **345**, 434
- [76] Thorne, K. S.: 1980, *Reviews of Modern Physics* **52**, 299
- [77] Thorne, K. S.: 1987, in S. W. Hawking and W. Israel (eds.), *300 Years of Gravitation*, Chapt. 9, p. 330, Cambridge University Press
- [78] Tinto, M., Estabrook, F. B., and Armstrong, J. W.: 2004, *Phys. Rev. D.* **69**, 082001
- [79] Tsubono, K.: 1995, in E. Coccia, G. Pizzella, and F. Ronga (eds.), *First Edoardo Amaldi Conference on Gravitational Wave Experiments*, p. 112
- [80] Unno, W., Osaki, Y., Ando, H., Saio, H., and Shibahashi, H.: 1989, *Non-radial Oscillations of Stars*, University of Tokyo Press, second edition
- [81] Ushomirsky, G., Cutler, C., and Bildsten, L.: 2000, *MNRAS* **319**, 902
- [82] Van Den Broeck, C.: 2005, *Class. Quant. Grav.* **22**, 1825
- [83] van der Klis, M., Wijnands, R. A. D., Horne, K., and Chen, W.: 1997, *ApJL* **481**, L97
- [84] Wagoner, R. V.: 1984, *ApJ* **278**, 345
- [85] Wagoner, R. V.: 2002, *ApJ* **578**, L63
- [86] Wainstein, L. and Zubakov, V.: 1962, *Extraction of Signals from Noise*, Prentice-Hall, Englewood Cliffs
- [87] Watts, A. L. and Strohmayer, T. E.: 2007, *Astrophys. & Space Science*, **308**, 625
- [88] Weisberg, J. M. and Taylor, J. H.: 1984, *Phys. Rev. Lett.* **52**, 1348
- [89] Willke, B., Aufmuth, P., Aulbert, C., Babak, S., et al.: 2002, *Class. Quant. Grav.* **19**, 1377
- [90] Zimmermann, M. and Szedenits, Jr., E.: 1979, *Phys. Rev. D.* **20**, 351