Robust Bayesian detection of unmodelled bursts

Antony C Searle

Centre for Gravitational Physics, Department of Physics, Faculty of Science, The Australian National University, Canberra ACT 0200 AUSTRALIA

E-mail: `antony.searle@anu.edu.au`

Abstract. A Bayesian treatment of the problem of detecting an unmodelled gravitational wave burst with a global network of gravitational wave observatories reveals that several previously proposed statistics have implicit biases that render them sub-optimal for realistic signal populations.

PACS numbers: 04.80.Nn,07.05.Kf,95.55.Ym

1. Introduction

Gürsel and Tinto [1] first noted that a system of *three* ground-based interferometers sensitive to *two* incoming gravitational wave polarisations over-determines the signal, permitting a consistency check in the absence of knowledge of the waveform. Their statistic was based on the failure to reject the null hypothesis that the observations were consistent with a gravitational wave from a postulated direction. Subsequent authors generalised the Gürsel-Tinto statistic to more than three interferometers and supplemented it with other consistency tests [2], and proposed various alternatives related to it [3, 4, 5, 6]. With the partial exception of [7], the use of Bayesian inference in the literature on coherent detection of unmodelled bursts has been confined to the justification of individual steps in the design of various statistics.

A fully Bayesian treatment of the problem requires (and uniquely follows from) the explicit specification of competing models of an experiment, including an explicit statement of how plausible we find any particular combination of the model parameters. This includes working out *a priori* how to 'spread our bets' over the large dimensional space of all possible signal waveforms. Such *prior plausibility distributions* are necessarily subjective—there is no 'right' nor 'wrong' way to choose them—but they do make definite statements about our expectations, and it is difficult to conceive of a scientist who would knowingly bet, for instance, that most detected gravitational waves will have a strain greater than unity.

The Bayesian requirement for a waveform prior stands in stark contrast to Gürsel-Tinto's advertised independence of signal waveform. The difference reflects a common criticism of Bayesian inference: that the use of priors introduces subjectivity and even bias into an analysis. One response to this criticism is to note that non-Bayesian statistics are not free of prior expectations about the world, but that their priors are merely implicit, unexamined and because of this may even contradict the stated intent of their designers. This suggests the question: what choice of priors will make a Bayesian statistic behave like the Gürsel-Tinto statistic and its relatives?

In this paper we set up a Bayesian analysis treating the experiment as a linear model with unspecified multivariate normal distributions as the priors on noise and signal. We then demonstrate how the Bayesian statistic reduces to (or limits to) Gürsel-Tinto and other related methods if we choose certain priors. The results are astonishing: none of the methods assume a uniform signal population across the sky, and most assume gravitational waves are unphysically large or undetectably small and very frequent. These priors are not 'wrong', but they are very far from the best guess of the scientists that designed them. Nor are the statistics necessarily ineffective, but they must be less effective than an practically implementable Bayesian statistic whose priors better reflect the real world.

2. Bayesian formulation

We can formulate much of the problem as a linear system. Pack N time series, each of an observatory's M measurements, into a vector

$$\mathbf{x} = [x_{1,1}, x_{1,2}, \ldots, x_{1,M}, x_{2,1}, \ldots, x_{N,M}]^T. \tag{1}$$

We could work equally well with any linear transformation of the time series (such as the Fourier or wavelet transforms).

One model for the noise in the system is that it is drawn from a multivariate normal distribution, described by an $MN \times MN$ covariance matrix $\mathbf{A}$. The plausibility of $\mathbf{x}$ assuming the noise-only hypothesis is then

$$p(\mathbf{x}|\text{`noise'}) = \frac{1}{(2\pi)^{MN/2}\sqrt{\det \mathbf{A}}} \exp -\frac{1}{2}\mathbf{x}^T\mathbf{A}^{-1}\mathbf{x}. \tag{2}$$

We may search for a *particular* signal

$$\mathbf{h} = [h_{+,1}, h_{+,2}, \ldots, h_{+,M}, h_{\times,1}, \ldots, h_{\times,M}] \tag{3}$$

which produces an additive response $\mathbf{Fh}$, where the $MN \times 2M$ response matrix $\mathbf{F}$ includes the antenna patterns and geometrical time delay associated with the *particular* direction (θ, ϕ) of the signal and the physical locations of the interferometers. The distribution of the data expected in the presence of this signal is

$$p(\mathbf{x}|\mathbf{h}) = \frac{1}{(2\pi)^{MN/2}\sqrt{\det \mathbf{A}}} \exp -\frac{1}{2}(\mathbf{x} - \mathbf{Fh})^T\mathbf{A}^{-1}(\mathbf{x} - \mathbf{Fh}). \tag{4}$$

This form is of no direct use because we do not know any *particular* signal $\mathbf{h}$ we want to search for. Bayesian marginalisation solves this problem for us but requires us to place a prior plausibility distribution on $\mathbf{h}$:

$$p(\mathbf{x}|\text{`signal'}) = \int_{\mathbf{h}} d\mathbf{h} \, p(\mathbf{h})p(\mathbf{x}|\mathbf{h}). \tag{5}$$

We must specify *a priori* how likely we think any particular signal is to occur. If the marginalisation integral is to be solved symbolically, this distribution must take the form of a multivariate normal distribution. However, some desirable signal models (those where the signal has fewer than $2M$ degrees of freedom) will result in singular covariance matrices. Instead we adopt a multivariate normal distribution with an $L \times L$ covariance matrix $\mathbf{Z}$ for abstract signal *parameters* $\mathbf{y}$ related to $\mathbf{h}$ through a $2M \times L$ matrix $\mathbf{W}$ of *template waveforms*, so that $\mathbf{h} = \mathbf{Wy}$.

A source with known waveforms $\mathbf{w}_+$ and $\mathbf{w}_\times$ whose relative amplitudes and projections onto the detector's polarisation basis is unknown (as is the case for a source of unknown polarisation angle and inclination) could be modelled with four amplitude parameters, $\mathbf{Z} = \mathbf{I}$ and

$$\mathbf{W} = \begin{bmatrix} \mathbf{w}_+ & \mathbf{w}_\times & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{w}_+ & \mathbf{w}_\times \end{bmatrix}. \tag{6}$$

This is a Bayesian analogue of a coherent matched filter analysis.

The *least informative* multivariate normal distribution prior has $2M$ parameters (one for each sample of each polarisation), with $\mathbf{Z} = \mathbf{I}$ and $\mathbf{W} = \sigma\mathbf{I}$ where σ is a scale factor. This prior says that we know nothing about h_i (save that it is unlikely to be very much larger that σ); we do not know if it is positive or negative or zero, and we do not know if it is correlated with h_j positively or negatively or not at all. This distribution also describes white noise, and we will call it a 'white noise' signal prior, even though it detects all signals, not just those we would classify as 'noise-like'.

Taking the number of parameters to be Y,

$$p(\mathbf{x}|\text{`signal'}) = \int_{\mathbf{y}} d\mathbf{y} \, p(\mathbf{y})p(\mathbf{x}|\mathbf{h}) \tag{7}$$

$$= \int_{\mathbf{y}} d\mathbf{y} \frac{\exp -\frac{1}{2}[(\mathbf{x} - \mathbf{FWy})^T\mathbf{A}^{-1}(\mathbf{x} - \mathbf{FWy}) + \mathbf{y}^T\mathbf{Z}^{-1}\mathbf{y}]}{(2\pi)^{(MN+Y)/2}\sqrt{\det \mathbf{A} \det \mathbf{Z}}}$$

$$= \frac{1}{(2\pi)^{MN/2}\sqrt{\det \mathbf{C}}} \exp -\frac{1}{2}\mathbf{x}^T \mathbf{C}^{-1}\mathbf{x} \qquad (8)$$

where $\mathbf{C}^{-1} = \mathbf{A}^{-1} - (\mathbf{A}^{-1}\mathbf{FW})[(\mathbf{FW})^T \mathbf{A}^{-1}(\mathbf{FW}) + \mathbf{Z}^{-1}]^{-1}(\mathbf{A}^{-1}\mathbf{FW})^T$.

This is just another multivariate normal distribution.

The Bayesian odds ratio tells us how likely the 'signal' hypothesis is, given the observation $\mathbf{x}$ in terms of how likely the observation is, given the two hypotheses under consideration and the prior plausibility of those hypotheses:

$$O = \frac{p(\text{`signal'}|\mathbf{x})}{p(\text{`noise'}|\mathbf{x})} = \frac{p(\text{`signal'})}{p(\text{`noise'})} \frac{p(\mathbf{x}|\text{`signal'})}{p(\mathbf{x}|\text{`noise'})}$$

$$= \frac{p(\text{`signal'})}{p(\text{`noise'})} \sqrt{\frac{\det \mathbf{A}}{\det \mathbf{C}}} \exp -\frac{1}{2}\mathbf{x}^T(\mathbf{C}^{-1} - \mathbf{A}^{-1})\mathbf{x}. \qquad (9)$$

We should choose the prior odds ratio $p(\text{`signal'})/p(\text{`noise'})$ to be $\ll 1$, reflecting the infrequency of detectable gravitational waves. (This is analogous to choosing the threshold λ for a frequentist likelihood.)

This is not yet a gravitational wave search; it looks for any gravitational wave of a known size from a known direction at a known time. In §4.2 we will see how to perform a Bayesian search. However, this odds ratio is the part of the Bayesian analysis that can be compared with previously proposed methods.

3. Comparison with existing methods

Several previously proposed tests have been expressed in the form

$$\mathbf{x}^T \mathbf{B}\, \mathbf{x} > \lambda \qquad (10)$$

for matrix $\mathbf{B}$ and threshold λ. The Bayesian test is instead

$$\frac{p(\text{`signal'})}{p(\text{`noise'})} \sqrt{\frac{\det \mathbf{A}}{\det \mathbf{C}}} \exp -\frac{1}{2}\mathbf{x}^T(\mathbf{C}^{-1} - \mathbf{A}^{-1})\mathbf{x} > 1 \qquad (11)$$

or, rearranging Equation 11,

$$\mathbf{x}^T(\mathbf{A}^{-1} - \mathbf{C}^{-1})\mathbf{x} > -2\ln \frac{p(\text{`signal'})}{p(\text{`noise'})} \sqrt{\frac{\det \mathbf{A}}{\det \mathbf{C}}} \qquad (12)$$

suggesting that by the appropriate choice of priors (including the signal model) we can create a Bayesian test that behaves like one of the previously proposed tests.

3.1. Tikhonov regularised statistic

Consider $\mathbf{A} = \mathbf{I}$, $\mathbf{W} = \sigma\mathbf{I}$ and $\mathbf{Z} = \mathbf{I}$ (white noise detectors and white noise signal of characteristic amplitude σ). Also let

$$\frac{p(\text{`signal'})}{p(\text{`noise'})} = e^{-\lambda/2}\sqrt{\det \sigma^2 \mathbf{F}^T\mathbf{F} + \mathbf{I}}. \qquad (13)$$

As $\mathbf{F}$ varies with direction, so too does the prior. For these priors, the Bayesian test of Equation 12 becomes

$$\mathbf{x}^T \mathbf{F}(\mathbf{F}^T\mathbf{F} + \sigma^{-2}\mathbf{I})^{-1}\mathbf{F}^T\mathbf{x} > \lambda \qquad (14)$$

which is the Tikhonov regularised statistic of [6] with regulariser $\alpha = \sigma^{-1}$.

The priors show that the Tikhonov regularised statistic implicitly assumes that gravitational wave bursts have a *particular size* (this physical interpretation of the regulariser was not made in [6]) and are *distributed non-uniformly on the sky* The strength of the directional bias depends on the size of the expected signal.

3.2. Gürsel-Tinto statistic

If we take the limit $\sigma \to \infty$ (the large signal limit) Equation 14 limits to

$$\mathbf{x}^T \mathbf{F} (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \mathbf{x} > \lambda \tag{15}$$

which is the Gürsel-Tinto statistic [1].

The priors show that the Gürsel-Tinto statistic implicitly assumes that gravitational wave bursts are *very large*, , *very frequent* (note that $p(\text{'signal'})/p(\text{'noise'}) \to \infty$), and *distributed non-uniformly on the sky*.

The effect of these priors is not so dramatic as we might expect; Gürsel-Tinto works in practice. Realistic signals are only an infinitesimal part of the signal population, but signals from that population occur infinitely frequently, and thus the method is able to detect them. Even so, one of Gürsel-Tinto's common failure modes is to misidentify a typical gravitational wave injection as a much larger (and therefore, it believes, much more likely to occur) gravitational wave with a different direction and polarisation that the network is nearly insensitive to.

3.3. Soft constraint statistic

Now allow the signal's characteristic amplitude to vary with direction, so that $\mathbf{W} = \sigma k(\Omega)\mathbf{I}$, and divide the Bayesian test in Equation 12 through by σ^2:

$$\frac{1}{\sigma^2} \mathbf{x}^T \mathbf{F} \left(\mathbf{F}^T \mathbf{F} + \frac{1}{\sigma^2 k^2(\Omega)} \mathbf{I} \right)^{-1} \mathbf{F}^T \mathbf{x} > -\frac{2}{\sigma^2} \ln \frac{p(\text{'signal'})/p(\text{'noise'})}{\sqrt{\det \sigma^2 k^2(\Omega) \mathbf{F}^T \mathbf{F} + \mathbf{I}}}.$$

Let

$$\frac{p(\text{'signal'})}{p(\text{'noise'})} = e^{-\lambda \sigma^2/2} \sqrt{\det \sigma^2 k^2(\Omega) \mathbf{F}^T \mathbf{F} + \mathbf{I}}. \tag{16}$$

Then in the limit of $\sigma \to 0$ we have

$$\frac{1}{k^2(\Omega)} \mathbf{x}^T \mathbf{F} \mathbf{F}^T \mathbf{x} > \lambda \tag{17}$$

which is the soft constraint statistic of [3] if the normalisation function $k^2(\Omega)$ is equal to the principal eigenvalue of $\mathbf{F}^T \mathbf{F}$ (though the normalisation function $k^2(\Omega) = \operatorname{tr} \mathbf{F}^T \mathbf{F}$ is required to give a direction-free null hypothesis distribution).

The priors show that the soft constraint statistic implicitly assumes gravitational wave bursts are *very small* and *very frequent* (note that $p(\text{'signal'})/p(\text{'noise'}) \to 1$). The Bayesian statistic can only draw very weak conclusions about the presence of very small signals (the division of Equation 12 by σ^2 is necessary to 'blow up' the its diminishing range), so even though the priors depend only weakly on direction they still strongly bias the soft constraint.

3.4. Hard constraint statistic

Now consider a linearly polarised signal model with polarisation angle some known function of direction $\psi(\Omega)$ and

$$\mathbf{W} = \sigma k(\Omega) \begin{bmatrix} \cos 2\psi(\Omega)\mathbf{I} \\ \sin 2\psi(\Omega)\mathbf{I} \end{bmatrix}. \tag{18}$$

Let

$$\frac{p(\text{'signal'})}{p(\text{'noise'})} = e^{-\lambda \sigma^2/2} \sqrt{\det \sigma^2 k^2(\Omega)(\mathbf{FW})^T \mathbf{FW} + \mathbf{I}}. \tag{19}$$
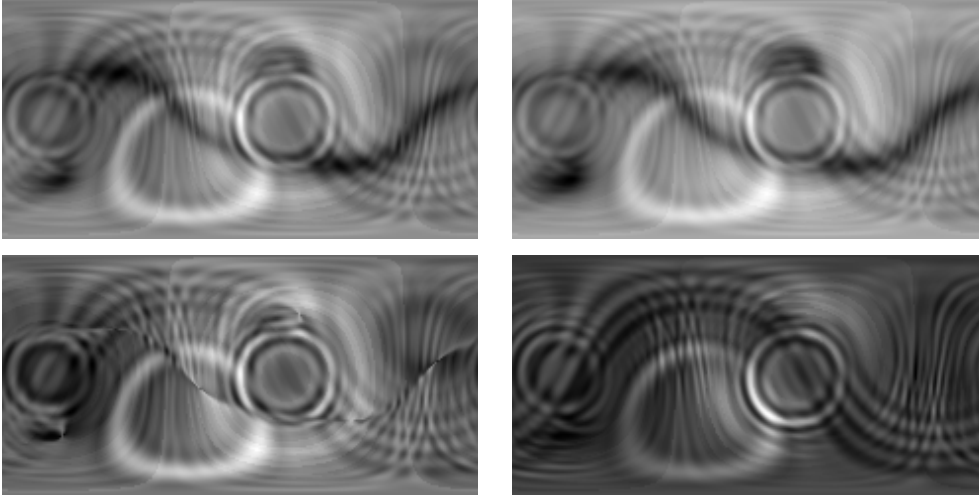
Figure 1. Four statistics as a function of direction for an injection of 1/16 seconds of white noise with amplitude SNR 5 into 3 identical white noise detectors sampled at 1024 Hz with the locations and orientations of LIGO Hanford, LIGO Livingston and Virgo. White is most plausible direction of the injection; the true direction corresponds to the centre of the maps. Top left: ln Bayesian odds ratio with $\sigma = 5$; the signal is localised with 95% confidence to 10^{-4} of the sky including the true direction. Top right: Tikhonov with $\alpha = 1/5$; this differs from the Bayesian case only by the direction prior. Bottom left: Gürsel-Tinto; note the discontinuities where the network becomes insensitive to one polarisation. Bottom right: Soft constraint with $k(\Omega) = \operatorname{tr} \mathbf{F}^T \mathbf{F}$. For this injection, the global maximums of the non-Bayesian statistics are not consistent with the true direction.

Then in the limit of $\sigma \to 0$ the Bayesian test Equation 12 divided by σ^2 goes to

$$\frac{1}{k^2(\Omega)} \mathbf{x}^T \mathbf{F} \mathbf{W} (\mathbf{F} \mathbf{W})^T \mathbf{x} > \lambda. \tag{20}$$

If the polarisation angle is defined as the one the network is most sensitive to, and the normalisation function is taken to be the principal eigenvalue of $(\mathbf{F}\mathbf{W})^T \mathbf{F}\mathbf{W}$, this expression is the hard constraint statistic of [3].

As well as the explicit design assumptions that gravitational wave bursts are *linearly polarised* and *optimally oriented*, the priors show that the hard constraint statistic implicitly assumes that gravitational wave bursts are *very small* and *very frequent* (note that $p(\text{'signal'})/p(\text{'noise'}) \to 1$), and (like the soft constraint) is strongly affected by the weak directional dependency of its priors.

3.5. Implications of implicit priors

The unphysical priors required to make a Bayesian analysis behave like any of the previously proposed methods indicates that none of these methods are optimal for the realistic scenarios their designers were targeting; they are optimal for scenarios greatly different, where gravitational waves are infinitely large or small and not uniformly distributed on the sky. A Bayesian analysis whose priors better reflect the true signal population will necessarily outperform all of these methods (see Figure 1), but the analysis above does not indicate how much better it will do.

What went wrong in the design of these methods? In the case of Gürsel-Tinto, the good idea of waveform independence was translated into an uninformative improper prior; but if an observation can be explained by either a moderately sized signal or an unphysically enormous signal, it is not optimal to treat these as equally credible alternatives. Ensuring that all directions have the same null hypothesis distribution counter-intuitively introduces a bias on the sky, because where the network is more sensitive the absence of excess power is actually stronger evidence against the signal hypothesis. The 'two detector paradox' described in [3] is simply an artefact of Gürsel-Tinto's uninformative prior. Of the *ad hoc* attempts to fix the 'paradox' described in [3, 6], only the Tikhonov regulariser constitutes an unambiguous theoretical improvement.

We have not considered here all the methods that have been proposed in the literature. Some of these, notably the other regularisation method in [6] and the almost-Tikhonov-regularisation in [5], likely suffer from similar problems.

4. Bayesian search

4.1. Robustness

All the methods discussed so far (including the Bayesian test) consider the only source of excess energy in interferometers to be a gravitational wave. Data orthogonal to span $\mathbf{F}$ is ignored because the noise and signal hypothesis make identical predictions and cannot be distinguished by it, even though that data is the Gürsel-Tinto null stream [1] and excess energy in it rejects the gravitational wave hypothesis. [2] supplemented Gürsel-Tinto with an *ad hoc* statistic in an attempt to solve this shortcoming.

A better approach is to include interferometer 'glitches' in the set of hypotheses under consideration; they are just another kind of signal model. A glitch can be represented as a different noise covariance distribution replacing $\mathbf{A}$. For example, if white noise glitches occur in white detectors with plausibility $p(\text{'glitch'})$, are additive and have characteristic amplitude γ, then we have a family of 2^N covariance distributions such as $\mathbf{A}_{01\ldots0} = \mathrm{diag}[\mathbf{I}, (1 + \gamma^2)\mathbf{I}, \ldots, \mathbf{I}]$ with prior plausibilities of the form $p(\text{'glitch'})^k[1 - p(\text{'glitch'})]^{2^N - k}$. The N-glitch hypothesis can 'explain' any data, but it incurs a large *Occam penalty* for its generality, whereas the signal hypothesis can more parsimoniously explain the tiny subspace of all possible data that is consistent with a gravitational wave. Signals from directions (or with a polarisation) that one detector is insensitive to are forced to compete against the less penalised $N - 1$ glitch hypotheses If we have more knowledge about glitches, we can use a set of glitch waveforms and parameters analogously to the construction of our signal hypothesis.

4.2. Marginalisation

To perform a search over parameters the model is not linear in, like time-of-arrival τ, sky direction Ω and signal model types, we need to numerically marginalise over these variables. For example, we can produce the odds ratio for a signal from *any* direction by marginalising over Ω:

$$O = \frac{p(\text{'signal'})}{p(\text{'noise'})} \frac{1}{p(\mathbf{x}|\text{'noise'})} \int_{S^2} \mathrm{d}\Omega \; p(\Omega)p(\mathbf{x}|\text{'signal'}). \qquad (21)$$

Previously proposed methods would instead *maximise* their statistic over their parameters, so we cannot directly compare this step, other than to note that for confidently detected signals the marginalisation will be dominated by a single peak and the integral will be proportional to the maximum of the integrand.

Signals from sources distributed uniformly in space will have a power-law distribution of characteristic amplitudes $p(\sigma) \propto \sigma^{-4}$ (this encodes the 'twice the sensitivity, eight times the event rate' rule of thumb). This provides us with a simple example of a parameterised set of signal models $\mathbf{W} = \sigma \mathbf{W}_1$ and a physically-inspired prior on the parameter. We may also want to look for signals parameterised in other ways. For example, we may want to consider different time and frequency bounds on time- and/or frequency-band-limited signals.

4.3. Implementation

If the noise and signal models are stationary, the Bayesian statistic for a particular signal size can be implemented (using a transformation the Fourier domain that diagonalises $\mathbf{C}$) as efficiently as the coloured noise versions of the previously proposed statistics considered here. If we declare the glitch and signal hypotheses exclusive (so that we will reject the very infrequent gravitational wave that occurs at the same time as a glitch) and model glitches in each detector as independent, we incur only a minimal extra cost to gain robustness against that population of instrumental artefacts. Marginalising over σ can reuse many intermediate values and requires relatively few samples for a steep σ^{-4} prior. Many other signal models present comparable opportunities for optimisation.

5. Conclusion

A simple Bayesian analysis of the problem of coherent detection of unmodelled gravitational wave bursts with a network of ground-based interferometric gravitational wave detectors is presented. It reduces or limits to several previously proposed methods for the same problem only when unphysical priors are used, indicating that those previously proposed methods cannot be optimal for realistic signal populations. A method for improving the robustness of the Bayesian method by substituting a more realistic noise model was noted, as were the steps for generalising to an all-sky unknown-arrival-time search. The Bayesian method can be implemented, depending on the signal model, to be computationally competitive with existing methods.

We can contrast the process of specifying physical models of the world, quantifying their predictions and forming the Bayesian statistic in §2 with the looser process of heuristically arguing our way to a statistic [1] and then noting its flaws and heuristically arguing our way to a mutated statistic [2, 3, 6]. In both cases, designers have freedom and subjective choices are made. In the Bayesian case, these are restricted to choosing the physical model and its priors; these have immediate physical interpretations which are explicit and can be contested. The Bayesian statistic follows uniquely from them. In the non-Bayesian case choices are made at every stage and their implications are uncertain; exactly what is the signal model after we have argued that we can safely disregard some of it [5]? Is a technique merely numerical or does it affect the physics of the system under consideration? The difference between the two ways of producing statistics can be neatly summed up as the difference between the statement 'we assume that gravitational waves have a characteristic size' and the

statement 'we apply Tikhonov regularisation to prevent rank deficiency in a matrix'. The opacity of the latter leaves us open to unexamined, unintended and unfortunate consequences.

References

[1] Y. Gursel and M. Tinto. *Phys. Rev. D*, 40:3884, 1989.
[2] S. Chatterji, A. Lazzarini, L. Stein, P. Sutton, A. Searle, and M. Tinto. Coherent network analysis technique for discriminating gravitational-wave bursts from instrumental noise. *Phys. Rev. D*, 74:082005, 2006.
[3] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher. *Phys. Rev. D*, 72:122002, 2005.
[4] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher. Constraint likelihood method: Generalization for colored noise. *J. Phys. Conf. Ser.*, 32:12, 2006.
[5] S. Mohanty, M. Rakhmanov, S. Klimenko, and G. Mitselmakher. Variability of signal to noise ratio and the network analysis of gravitational wave burst signals. *Class. Quant. Grav.*, 23:4799–4809, 2006.
[6] M. Rakhmanov. Rank deficiency and tikhonov regularization in the inverse problem for gravitational-wave bursts. *Class. Quant. Grav.*, 23:S673–S685, 2006.
[7] W. G. Anderson, P. R. Brady, J. D. E. Creighton, and É. É. Flanagan. *Phys. Rev. D*, 63:042003, 2001.