LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Document Type	**LIGO-T030041-01-Z**	2003/11/24

Notes on bounding the binary inspiral rate

Lee Samuel Finn

Distribution of this draft:

LIGO I Collaboration

California Institute of Technology
LIGO Project - MS 51-33
Pasadena CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project - MS 20B-145
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

WWW: http://www.ligo.caltech.edu/

Processed with LaTeX on 2003/11/24

Abstract

The initial goal of the LIGO search for inspiraling compact binaries is to bound their rate. Candidate gravitational wave events generated by the existing binary inspiral search are characterized by a signal amplitude, a chirp mass, and a characteristic event time. The source population, from which true binary inspiral gravitational wave events are drawn, has binaries distributed with galaxies and, within our own galaxy, following a distribution not too different than that associated with pulsars. In the S1 analysis the totality of candidate events are characterized by the largest amplitude event observed. In this note we show how to make fuller use of the number of events and their amplitude in bounding the binary inspiral event rate. The analysis described here is an application of an analysis first described in [1].

$Id: T030041.tex,v 1.7 2003/03/24 01:59:39 lsf Exp $

Contents

1 Introduction

The initial goal of the LIGO search for inspiraling compact binaries is to bound their rate. Candidate gravitational wave events generated by the existing binary inspiral search are characterized by a signal amplitude, a chirp mass, and a characteristic event time. The source population, from which true binary inspiral gravitational wave events are drawn, has binaries distributed with galaxies and, within our own galaxy, following a distribution not too different than that associated with pulsars. In the S1 analysis the totality of candidate events are characterized by the largest amplitude event observed. In this note we show how to make fuller use of the number of events and their amplitude in bounding the binary inspiral event rate. The analysis described here is an application of an analysis first described in [1].

2 Source population and gravitational wave strain events

In this section we define the source population and relate it to the population's contribution to the events identified in the analysis.

Begin by characterizing the *local specific rate density* of binary coalescence $\mathcal{N}$:

$$\dot{\mathcal{N}} \equiv \frac{d^5 N}{dt\, d^3x\, d\mathcal{M}}, \tag{1}$$

where d^3x is a spatial volume element and $\mathcal{M}$ is the chirp mass that charactrerizes each binary. We assume that the local specific rate density is a function of location $\vec{r}$ and chirp mass $\mathcal{M}$. The local specific rate density fully characterizes the source population.

It is convenient to "factor" $\dot{\mathcal{N}}$ into several components, each with a physical meaning:

$$\dot{\mathcal{N}} = \dot{n}_S f(\vec{x}|\mathcal{I}) P(\mathcal{M}|\vec{x}, \mathcal{I}) \tag{2}$$

where

$$\dot{n}_S = \left(\text{total rate of binary coalescence in a "standard" galaxy} \right), \tag{3}$$

$$P(\mathcal{M}|\vec{x}, \mathcal{I}) = \left(\begin{array}{l} \text{probability that a binary coalescence at } \vec{x} \text{ (in} \\ \text{the galaxy located at } \vec{x}\text{) has chirp mass } \mathcal{M} \end{array} \right), \tag{4}$$

$$f(\vec{x}|\mathcal{I}) = \left(\begin{array}{l} \text{Binary coalescence rate at } \vec{x} \text{ relative to the rate in a standard} \\ \text{galaxy; a density with units [standard galaxies/volume]} \end{array} \right), \tag{5}$$

$$\mathcal{I} = \left(\begin{array}{l} \text{parameters describing the coalescing binary} \\ \text{distribution model in space and chirp mass} \end{array} \right). \tag{6}$$

Note that $P(\mathcal{M}|\vec{r}, \mathcal{I})$ may be independent of $\vec{r}$. For the purpose of binary inspiral observations $\mathcal{I}$ might represent our parameterization of the galactic source distribution (density in and size of bulge, scale height, etc.), the variations among galaxy type (density in ellipticals vs. spirals), and the variations in the distribution of binaries by chirp mass (which may also depend on galaxy type). Here and henceforth we assume that $\mathcal{I}$ is fixed and consider how gravitational wave observations can bound $\dot{n}_S$, which we take to be the rate in our own galaxy. Having said that, however, note that nothing we do here or below excludes the possibility of using the observations to bound additional model parameters included in $\mathcal{I}$.

The population model $\mathcal{I}$ leads to a distribution of gravitational wave events incident on the detector array. Each individual event in this distribution is characterized by two wave polarization amplitudes $(h_+, h_\times)$, a chirp mass $\mathcal{M}$ and a wave propagation direction $\vec{n}$.[1] Denote this parameterization of a real source event by $\vec{h}$,

$$\vec{h} = \left(\begin{array}{l} \text{Parameterization of binary inspiral} \\ \text{wave event incident on detector array} \end{array} \right), \tag{7}$$

and write the distribution of real source events in the population model $\mathcal{I}$ by

$$p(\vec{h}|\mathcal{I}) = \left(\begin{array}{l} \text{probability of strain event characterized} \\ \text{by } \vec{h} \text{ given population characterized by } \mathcal{I} \end{array} \right). \tag{8}$$

[1] The polarization amplitudes measured at the detectors depend on the internal orientation of the binary system's angular momentum axis relative to the propagation direction of the radiation. The angular momentum points in a random direction and the effect this has on the distribution of h_+ and $h_\times$ is readily calculated: see, e.g., [2].

An obvious "Olber's Paradox" afflicts the definition of $p(\vec{h}|\mathcal{I})$. Distant sources are more numerous than closer ones; correspondingly, weaker sources are more numerous than stronger ones. Without a source strength cut-off the fraction of sources with non-zero intensity is zero. In the real universe the paradox is resolved in the usual way by the cosmological redshift and the finite age of the universe. Intuitively we know that the details of the cut-off are not important as long as the events with amplitude less than the cut-off do not contribute to the events that are actually observed; we will see in quantitative detail how this comes to pass below.

A particular inspiral event $\vec{h}$ incident on the detector may or may not lead to an observed event. Denote the data processing pipeline that leads to the identification of gravitational wave event candidates by

$$\mathcal{J} = \left(\begin{array}{l}\text{description of the analysis}\\ \text{method that identifies}\\ \text{gravitational wave events}\end{array}\right) \tag{9}$$

and write the probability that the real gravitational wave event characerized by $\vec{h}$ leads to a candidate event by $\epsilon(\vec{h},\mathcal{J})$:

$$\epsilon(\vec{h},\mathcal{J}) = \left(\begin{array}{l}\text{probability that the event } \vec{h}\\ \text{gives rise to a detector event}\\ \text{when analyzed according to } \mathcal{J}\end{array}\right) \tag{10}$$

Note that $\epsilon(\vec{h},\mathcal{J})$ depends on the analysis that identifies events (here represented by $\mathcal{J}$), the noise character and the instrument calibration.[2] With these specified, however, $\epsilon(\vec{h},\mathcal{J})$ is readily determined by simulation.

Each *observed* event is characterized by a set of parameters that we denote $\vec{H}$:

$$\vec{H} = \left(\begin{array}{l}\text{Parameters describing gravitational}\\ \text{wave event identified in detector array}\end{array}\right). \tag{11}$$

At the very least $\vec{H}$ will include some measure of the event amplitude and chirp mass. It may also include a locus of points denoting the possible origin of the incident gravitational wave on the sky. With enough detectors and a sufficiently sophisticated analysis it may be that $\vec{H}$ includes the amplitudes in both polarizations, a single location on the sky, and other information about the character of the burst; however, for the remainder of this note we assume that $\vec{H}$ represents estimated signal-to-noise ratio and chirp mass. As above, however, nothing we do here or below excludes the possibility that $\vec{H}$ is of higher dimension.

A critical relationship is the one between actual events, described by $\vec{h}$, and detected events, described by $\vec{H}$. Let $q(\vec{H}|\vec{h})$ be the probability that the real event $\vec{h}$, *if observed*, leads to the characterized observation $\vec{H}$:

$$q(\vec{H}|\vec{h},\mathcal{J}) = \left(\begin{array}{l}\text{probability that } \vec{H} \text{ characterizes}\\ \text{the observed event associated}\\ \text{with actual event } \vec{h} \text{ and}\\ \text{identified by method } \mathcal{J}\end{array}\right). \tag{12}$$

[2]We assume here that the noise and instrument calibration are stationary. Analysis in the presence of non-stationary noise and/or instrument calibration is described in section 7.1.

The probability $q(\vec{H}|\vec{h},\mathcal{J})$ can be thought of as the uncertainty in the determination of the character of the signal described by $\vec{h}$; alternatively, it can be thought of as the "point spread function" for binary inspiral observations. Note that $q(\vec{H}|\vec{h})$ depends on the nature of the analysis $\mathcal{J}$, the nature of the detector noise, and the calibration. Once these are specified, however, $q(\vec{H}|\vec{h})$, like $\epsilon(\vec{h},\mathcal{J})$, is readily determined by simulation.

The contribution to the detector output of gravitational wave events associated with the population $\mathcal{I}$ is thus described by the *foreground event distribution*

$$
\begin{aligned}
P_F(\vec{H}|\mathcal{I}\mathcal{J}) &= \left(\text{ probability of making the observation} \vec{H} \text{ given the population } \mathcal{I}\ \right) && (13)\\
&\propto \int d^n h\, q(\vec{H}|\vec{h},\mathcal{J})\epsilon(\vec{h},\mathcal{J})p(\vec{h}|\mathcal{I}) && (14)
\end{aligned}
$$

where $d^n h$ is the measure on $\vec{h}$: i.e., $dh_+\, dh_\times\, d\mathcal{M}, d^2S$, where d^2S is the surface element on the sphere described by orientation of the binary relative to the wave propagation direction. Note that $P_F(\vec{H}|\mathcal{I}\mathcal{J})$ is completely described by the population model $p(\vec{h}|\mathcal{I})$ and the analysis pipeline $\mathcal{J}$, as characterized by simulation.

We now see the conditions under which we can impose a cut-off on the source population model: i.e., how we can ignore sources whose signal-to-noise is sufficiently weak in constructing our source population model $\mathcal{I}$ or $p(\vec{h}|\mathcal{I})$. Note that $P_F(\vec{H}|\mathcal{I}\mathcal{J})$ is independent of

- the absolute normalization of $p(\vec{h}|\mathcal{I})$;
- the behavior of $p(\vec{h}|\mathcal{I})$ for events $\vec{h}$ for which $\epsilon(\vec{h},\mathcal{J})$ is negligible.

Consequently, in defining our population model $\mathcal{I}$ (and, thus, $p(\vec{h}|\mathcal{I})$) we can assume that there are no events whose signal strength at the detector is so low that the probability of detection $\epsilon(\vec{h},\mathcal{J})$ is negligible.

Associated with the source population is the standard galaxy event rate $\dot{n}_S$. Not every source event leads to an observed event. The fraction of source events that lead to observed events is the total detection efficiency, which depends on the source population model, the detector noise and calibration, and the analysis methodology that identifies gravitational wave events. Writing the total efficiency as $\epsilon(\mathcal{I}\mathcal{J})$ we have[3]

$$
\begin{aligned}
\dot{n}_F &= \left(\text{ rate of observed foreground events }\right) && (15)\\
&= \int d^3x d\mathcal{M}\dot{\mathcal{N}} && (16)\\
&= \dot{n}_S f(\mathcal{I})\epsilon(\mathcal{I}\mathcal{J}) && (17)
\end{aligned}
$$

where

$$
\begin{aligned}
\epsilon(\mathcal{I}\mathcal{J}) &= \int d^n h\, \epsilon(\vec{h},\mathcal{J})p(\vec{h}|\mathcal{I}) && (18)\\
f(\mathcal{I}) &= \int d^3x\, f(\vec{x}|\mathcal{I}). && (19)
\end{aligned}
$$

[3]Again, we assume here stationary detector noise, calibration, etc., and treat the case of non-stationarity in section 7.1 below.

3 Background distribution and event rate

Observed events may arise from the source population, in which case they are drawn from the distribution $P_F(\vec{H}|\mathcal{I}\mathcal{J})$, or from environmental or instrumental artifacts. We refer to the distribution of events associated with environmental or instrumental artifacts as the background distribution:

$$P_B(\vec{H})d^nH = \begin{pmatrix} \text{fraction of background events in} \\ \text{the } n\text{-ball of radius } d^nH \text{ about } \vec{H} \end{pmatrix}. \tag{20}$$

Background events occur at a given rate, which we denote $\dot{n}_B$:

$$\dot{n}_B = \begin{pmatrix} \text{rate of background events} \\ \text{of any amplitude, chirp mass} \end{pmatrix} \tag{21}$$

Both the background distribution and its rate $\dot{n}_B$ may be estimated from time-delay coincidence analysis assuming that there is no preference for “zero-delay” background disturbances in the gravitational wave channel that cannot be vetoed by other means.

4 The likelihood function

The likelihood is the probability of a particular observation under a fixed hypothesis. In our case the hypothesis is that there is a source population $\mathcal{I}$ characterized by the standard galaxy rate $\dot{n}_S$ and our observation is a set of N observed events $\vec{H}_k$:

$$\mathcal{H} = \left\{\vec{H}_n : n = 1 \ldots N\right\}. \tag{22}$$

Focus first on the probability of a single event $\vec{H}$. That event may be foreground or background. The rate of foreground events is the product of the total detection efficiency $\epsilon(\mathcal{I}\mathcal{J})$ (cf. eqn. 18) and the standard galaxy rate $\dot{n}_S$, which is what we wish to determine. Write the foreground event rate in terms of the background event rate $\dot{n}_B$ and a parameter α, $\alpha \in [0, 1)$:

$$\begin{aligned} \dot{n}_F &= f(\mathcal{I})\epsilon(\mathcal{I}\mathcal{J})\dot{n}_S & (23) \\ &= \dot{n}_B \frac{\alpha}{1-\alpha}. & (24) \end{aligned}$$

As defined the parameter α is the probability that a particular event is a foreground event. In terms of α the probability of a particular event $\vec{H}$ is thus

$$P(\vec{H}|\mathcal{I}\mathcal{J}, \dot{n}_B, \dot{n}_S) = (1-\alpha)\, P_B(\vec{H}) + \alpha P_F(\vec{H}|\mathcal{I}\mathcal{J}). \tag{25}$$

Now assume that gravitational wave events are independent of each other, and that the same is true of background events. The probability of making the particular observation $\mathcal{H}$ is then product of the probability of observing N events, which is given by the Poisson distribution, and the probability that the N observed events are characterized by their particular $\vec{H}_k$, or

$$P(\mathcal{H}|T, \dot{n}_B, \dot{n}_S, \mathcal{I}\mathcal{J}) = P(N|\mu) \begin{cases} 1 & N = 0 \\ \prod_{k=1}^{N} P\left(\vec{H}_k|\dot{n}_B, \dot{n}_S, \mathcal{I}\mathcal{J}\right) & N > 1 \end{cases} \tag{26}$$

where

$$P(N|\mu) = \frac{\mu^N}{N!} e^{-\mu} \tag{27}$$

is the Poisson distribution and

$$\mu = T\left[\dot{n}_B + \dot{n}_S f(\mathcal{I})\epsilon(\mathcal{I}\mathcal{J})\right] \tag{28}$$

is the expected number of events in an observation of livetime T.

From the likelihood and the observation $\mathcal{H}$ we can find the bounds on $(h_0, \dot{n}_S)$ using Bayesian techniques. In the next section we describe a Frenquentist analysis for $(h_0, \dot{n}_S)$.

5 A Frequentist analysis

Having determined the likelihood $P(\vec{H}|\mathcal{I}\mathcal{J}, \dot{n}_B, \dot{n}_S)$, the probability of observing the single event $\vec{H}$ under the hypothesis $(\mathcal{I}\mathcal{J}, \dot{n}_B, \dot{n}_S)$, we can proceed to find from $\mathcal{H}$ the bound on $\dot{n}_S$ (and other parameters that characterize the population model $\mathcal{I}$) using a Frequentist analysis.

Begin by introducing a partition of the parameter space spanned by $\vec{H}$. For example,

- if $\vec{H}$ is just the event amplitude, then introduce a partition in amplitude;
- if $\vec{H}$ is the amplitude in two different polarizations and the wave propagation direction, then introducing a partition in the amplitude in each polarization and a partition on the sphere for the wave propagation direction.

Bin the individual events in the observation $\mathcal{H}$ according to this partition of the parameter space.

From our knowledge of $P(\vec{H}|\mathcal{I}\mathcal{J}, \dot{n}_B, \dot{n}_S)$ and the observation duration T, and under the assumption that the events are independent, we know the expected number of events in each bin introduced above. Form $\chi^2(\mathcal{H}|\mathcal{I}\mathcal{J}, \dot{n}_B, \dot{n}_S, T)$, the χ^2 statistic for the observation $\mathcal{H}$ as binned. Note that, because we expect a particular number of events (we have specified T, $\dot{n}_B$ and $\dot{n}_S$ and $\mathcal{J}$) the χ^2 statistic is drawn from the χ^2 distribution with N degrees of freedom, where N is the number of single events in $\mathcal{H}$.

We can now ask what hypotheses $\dot{n}_S$ (or $\dot{n}_S$ and $\mathcal{I}$) lead to $\chi^2(\mathcal{H}|\mathcal{I}\mathcal{J}, \dot{n}_B, \dot{n}_S, T)$ such that the probability of obtaining this χ^2 is greater than, e.g., 90%. While not a confidence interval in the usual sense, it has a similar interpretation as the range of hypotheses for which the observation is likely, in the χ^2 sense.

This approach has the added advantage that it provides, coincidentally, a measure of goodness-of-fit: in particular, if there is no set of hypotheses that include the observation as likely, no “confidence interval” will be returned.

6 Relation to the “most luminous event” analysis

In the S1 analysis the set of N observations $\mathcal{H}$ is characterized by the single observation $\vec{H}_0 \in \mathcal{H}$ with the largest signal-to-noise, ρ_0. It is certainly the case that all foreground events, regardless of their number, have signal-to-noise less than ρ_0. Given a population model $\mathcal{I}$ and a method for identifying events $\mathcal{J}$ the fraction of foreground events whose signal-to-noise is less than or equal to ρ_0 is

$$C_F(\rho_0|\mathcal{I}\mathcal{J}) = \int_0^{\rho_0} d\rho \int_0^{\infty} d\mathcal{M}\, P_F(\rho, \mathcal{M}|\mathcal{I}\mathcal{J}), \tag{29}$$

where we have assumed that $\vec{H}_0$ is completely determined by the estimated chirp mass $\mathcal{M}_0$ and signal-to-noise ρ_0 (cf. eq. 14). The probability that all foreground events, regardless of their number, have signal-to-noise ρ_F less than ρ_0 is thus

$$\begin{aligned} P(\rho_F < \rho_0|T, \dot{n}_S, \mathcal{I}\mathcal{J}) &= \sum_{n=0}^{\infty} P(N|T\dot{n}_F) C_F(\rho_0|\mathcal{I}\mathcal{J})^n & (30) \\ &= e^{-T\dot{n}_F[1-C_F(\rho_0|\mathcal{I}\mathcal{J})]} & (31) \end{aligned}$$

(where the absence of any foreground events of course means that no foreground events have signal-to-noise greater than ρ_0).

Associated with this probability is the probability density that the most luminous event has signal-to-noise ρ_0:

$$\begin{aligned} p_F(\rho_F|T, \dot{n}_F, \mathcal{I}\mathcal{J}) &= \frac{d}{d\rho_0} e^{-T\dot{n}_F[1-C_F(\rho_0|\mathcal{I}\mathcal{J})]} & (32) \\ &= T\dot{n}_S P_F(\rho_0|\mathcal{I}\mathcal{J}) e^{-T\dot{n}_F[1-C_F(\rho_0|\mathcal{I}\mathcal{J})]}. & (33) \end{aligned}$$

This probability density is also the likelihood for the observation that the most luminous foreground event has signal-to-noise ρ_0. In the most luminous event analysis one assumes that the most luminous event is a foreground event and then, from this likelihood, uses a Bayesian analysis with a uniform prior in $\dot{n}_S$ to determine a credible set, bounded below by $\dot{n}_S = 0$, associated with a probabiltiy p. The upper end of this credible set is taken to be the upper limit on $\dot{n}_S$. (See [3] for more details and an alternative interpretation of the most luminous event.)

The critical difference between the “most luminous event” analysis and the more comprehensive analysis in the previous sections is that “most luminous event” analyses discard virtually all the information gathered in an observation: the number of events observed, the distribution of events in amplitude, etc.

7 Discussion

7.1 Non-stationarity

The detector noise and calibration are not steady over the entire observation T. We can accommodate a time-varying noise and calibration if we can treat the noise and calibration as piecewise constant in time and know in what interval each of the N events in the observation $\mathcal{H}$ occurs.

Partition the total observation time T into M sub-intervals of duration t_k, $\sum_k^M t_k = T$, in which the noise and calibration are constant. Similarly partition the observation $\mathcal{H}$ into M disjoint sub-observations $\mathcal{H}_k$, with the union of the $\mathcal{H}_k$ equal to $\mathcal{H}$, such that all the events in $\mathcal{H}_k$ occur in the interval t_k. Associated with each sub-observation is the likelihood of making that observation given the expected background rate in the given interval: $P\left(\mathcal{H}_k|t_k, \dot{n}_{B,k}, \dot{n}_S, \mathcal{I}\mathcal{J}\right)$. *Note that the* *background event rate $\dot{n}_B$ and the distributions $P_B(\vec{H})$ and $P_F(\vec{H})$ will in general be different in each sub-interval.* The likelihood for the complete observation of duration T is then

$$P\left(\mathcal{H}\middle|\left\{t_k, \dot{n}_{B,k}\right\}, \dot{n}_S, \mathcal{I}\mathcal{J}\right) = \prod_{k=1}^{M} P\left(\mathcal{H}_k | t_k, \dot{n}_{B,k}, \dot{n}_S, \mathcal{I}\mathcal{J}\right). \tag{34}$$

From the likelihood we can derive the bound on $\dot{n}_S$ in the usual way.

Handling non-stationarity thus reduces to identifying epochs over which the noise and calibration are approximately stationary. Residual non-stationarity in each epoch will lead to a systematic error in the analysis. The degree to which stationarity should be required in an epoch is thus set by the level of the other systematic errors in the analysis.

Tracking calibration line amplitudes provides one method of identifying epochs over which the calibration is stationary. Observing the time dependent rate of background events and using a Bayesian Block analysis [4] is a possible approach to determining epochs when the noise is stationary.

7.2 Background rate uncertainty

The background rate $\dot{n}_B$ is determined experimentally. Associated with the experimental background rate is an uncertainty. Let

$$P_B\left(\dot{n}_B\right) d\dot{n}_B = \left(\text{ degree of belief that } \dot{n}_B \text{ is in } [\dot{n}_B, \dot{n}_B + d\dot{n}_B).\ \right) \tag{35}$$

We can marginalize the likelihood over this uncertainty, obtaining a new likelihood that is independent of uncertain $\dot{n}_B$

$$P(\mathcal{H}|T, \dot{n}_S, \mathcal{I}\mathcal{J}) = \int d\dot{n}_B P_B(\dot{n}_B) P(\mathcal{H}|\dot{n}_B, \dot{n}_S, \mathcal{I}\mathcal{J}) \tag{36}$$

The uncertainty $P_B\left(\dot{n}_B\right)$ may be estimated by making many estimates of the background rate, all at different delays, as long as the delays are much greater than any residual correlation time in the input time series from which the events are determined.

7.3 Goodness-of-fit

Any observation $\mathcal{H}$ will yield a bound on $\dot{n}_S$, even if the observation is, itself, very unlikely given our state of knowledge regarding the expected distribution of events in $\vec{H}$, background rate $\dot{n}_B$, and source population model $\mathcal{I}$. Given our best estimate of $\dot{n}_S$ we can ask whether the corresponding model is consistent with the observations using a χ^2 test.

Alternatively, the value of the likelihood for the observation $\mathcal{H}$ provides a measure of the degree to which the observations are expected in the context of the model. Focus attention on the maximum of the likelihood over the source rate $\dot{n}_S$ given the observation. Simulations for this $\dot{n}_S$ will determine a distribution of observations and, correspondingly, values of the likelihood under the assumption that the rate is $\dot{n}_S$. The value of the likelihood for the actual observation can be compared to this distribution in order to determine how exceptional the observation is. If the observation is too exceptional given the best-fit (i.e., the maximum likelihood value of) $\dot{n}_S$ then we may wish to regard the bound on $\dot{n}_S$ as suspect.

8 References

[1] Lee Samuel Finn. Notes on rate vs. strength upper limits. Technical Report T030017, Laser Interferometer Gravitational Wave Observatory, LIGO Document Control Center, California Institute of Technology, 2003.

[2] Lee Samuel Finn. Binary inspiral, gravitational radiation, and cosmology. *Phys. Rev. D*, 53(6):2878–2894, 15 March 1996.

[3] Lee Samuel Finn. Comments on inspiral event rate upper limit determination. Technical Report LIGO-T030048-Z, Laser Intereferometer Gravitational Wave Observatory, March 2003. Internal working note.

[4] J. D. Scargle. Studies in Astronomical Time Series Analysis. V. Bayesian Blocks, a New Method to Analyze Structure in Photon Counting Data. *Astrophys. J.*, 504:405–+, September 1998.