Draft version April 7, 2009
Preprint typeset using LaTeX style emulateapj v. 10/09/06

EXPLORING SHORT GAMMA-RAY BURSTS AS GRAVITATIONAL-WAVE STANDARD SIRENS

Samaya Nissanke[1,2], Scott A. Hughes[2], Daniel E. Holz[3], Neal Dalal[1], Jonathan L. Sievers[1]

Draft version April 7, 2009

ABSTRACT

Recent observations support the hypothesis that a large fraction of "short-hard" gamma-ray bursts (SHBs) are associated with the inspiral and merger of compact binaries. Since gravitational-wave (GW) measurements of well-localized inspiraling binaries can measure absolute source distances with high accuracy, simultaneous observation of a binary's GWs and SHB would allow us to directly and independently determine both the binary's luminosity distance and its redshift. Such a "standard siren" (the GW analog of a standard candle) would provide an excellent probe of the relatively nearby ($z \lesssim 0.3$) universe's expansion, independent of the cosmological distance ladder, and thus complementing other standard candles. Previous work explored this idea using a simplified formalism to study measurement by advanced GW detector networks, incorporating a high signal-to-noise ratio limit to describe the probability distribution for measured parameters. In this paper we eliminate this simplification, constructing distributions with a Markov Chain Monte Carlo technique. We assume that each SHB observation gives both the source sky position and the time of coalescence, and we take both binary neutron stars and black hole-neutron star coalescences as plausible SHB progenitors. We examine how well parameters (particularly the luminosity distance) can be measured from GW observatations of these sources by a range of ground-based detector networks. We find that earlier estimates overstate how well distances can be measured, even at fairly large signal-to-noise ratio. The fundamental limitation to determining distance to these sources proves to be the gravitational waveform's degeneracy between luminosity distance and source inclination. Despite this, we find that excellent results can be achieved by measuring a large number of coalescing binaries, especially if the worldwide network consists of many widely separated detectors. Advanced GW detectors will be able to determine the absolute luminosity distance to an accuracy of 10–30% for NS-NS and NS-BH binaries out to 600 and 1400 Mpc, respectively.

Subject headings: cosmology: distance scale—cosmology: theory—gamma rays: bursts—gravitational waves

1. INTRODUCTION

1.1. *Overview*

There are presently two operational multikilometer interferometric gravitational-wave (GW) detectors: LIGO[4] and Virgo[5]. They are sensitive to the GWs produced by the coalescence of two neutron stars to a distance of roughly 30 Mpc, and to the coalescence of a neutron star with a $10M_{\odot}$ black hole to roughly 60 Mpc. Over the next several years these detectors will undergo upgrades which are expected to extend their range by a factor ~ 10. At these advanced sensitivity levels, most estimates suggest that detectors should measure at least a few, and possibly a few dozen, binary coalescence events every year (e.g., Kopparapu et al. 2008).

It has long been argued that neutron star-neutron star (NS-NS) and neutron star-black hole (NS-BH) mergers are likely to be accompanied by a gamma-ray burst (Eichler et al. 1989). Recent evidence supports the hypothesis that many short-hard gamma-ray bursts (SHBs) are indeed associated with such mergers (Fox et al. 2005, Nakar et al. 2006, Berger et al. 2007, Perley et al. 2008). This suggests the exciting possibility that it may be possible to simultaneously measure a binary coalescence in gamma rays (and associated afterglow emission) and in GWs. The combined electromagnetic and gravitational view of these objects is likely to teach us substantially more than what we learn from either data channel alone. Because GWs track a system's global mass and energy dynamics, measuring GWs from a coalescing binary allows us to determine with exquisite accuracy "intrinsic" binary properties, such as the masses and spins of its members. As we describe in the following subsection, GWs can also determine a system's "extrinsic" properties, such as location on the sky and distance to the source. In particular, the amplitude of a binary's GWs directly encodes its luminosity distance. Direct measurement of a coalescing binary could thus be used as a cosmic distance measure: Binary inspiral would be a "standard siren" (the GW equivalent of a standard candle, so-called due to the sound-like nature of GWs) whose calibration depends only on the validity of general relativity (Dalal et al. 2006).

Unfortunately, GWs alone do not measure extrinsic parameters as accurately as the intrinsic ones. As we describe in more detail in the following section, in general a GW observation of a binary measures a complicated combination of the distance to the binary, the binary's position on the sky, and the binary's orientation, with overall fractional accuracy $\sim 1/\text{signal-to-noise}$. As the distance is degenerate with the angular parameters, us-

[1] CITA, University of Toronto, 60 St. George St., Toronto, ON, M5S 3H8, Canada
[2] Department of Physics and MIT Kavli Institute, 77 Massachusetts Avenue, Cambridge, MA 02139
[3] Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545
[4] http://www.ligo.caltech.edu
[5] http://www.virgo.infn.it

ing GWs to measure absolute distance to a source requires a mechanism to break the degeneracy. Associating the GW coalescence waves with a short-hard gamma-ray burst (SHB) is a near-perfect way to break this degeneracy. In this paper we explore the ability of near-future GW detector networks (such as LIGO and Virgo) to constrain binary parameters (and in particular, distance to a binary), when used in conjunction with electromagnetic observations of the same event (such as an associated SHB). We also examine how well these measurements can be improved if planned detectors in Western Australia (AIGO[6]) and in Japan's Kamioka mine (LCGT[7]) are operational. This paper substantially updates and improves upon earlier work (Dalal et al. 2006, hereafter DHHJ06), utilizing a significantly more sophisticated and accurate parameter estimation technique. In the next section we review standard sirens, while in Sec. 1.3 we briefly summarize DHHJ06. We follow this with a subsection describing the organization and background relevant for the rest of the paper.

1.2. *Standard sirens*

The GWs produced by the inspiral of two compact bodies directly encode the luminosity distance to the binary. It has long been recognized that GW inspiral measurements could thus be used as powerful tools for cosmology. Schutz (1986) first demonstrated this by analyzing how binary coalescences allow a direct measurement of the Hubble constant; Marković (1993) and Finn & Chernoff (1993) subsequently generalized this approach to include other cosmological parameters. More recently, there has been much interest in the measurements enabled when the GWs from a merger are accompanied by a counterpart in the electromagnetic spectrum (Bloom et al. 2009, Phinney 2009, Kulkarni & Kasliwal 2009). In this paper we focus exclusively on GW observations of binaries that have an independent sky position furnished by electromagnetic observations (e.g., of gamma-rays associated with an SHB, or an accompanying optical afterglow).

We begin by examining the gravitational waves from a binary inspiral as measured in a single GW detector. For simplicity, we only present here the lowest order contribution to the waves; in our subsequent calculations, our results are taken to higher order (see Sec. 2.1). The waveform generated by a source at luminosity distance D_L, corresponding to redshift z, is given by

$$h_+ = \frac{2(1+z)\mathcal{M}}{D_L}\left[\pi(1+z)\mathcal{M}f\right]^{2/3}\left(1+\cos^2\iota\right)\times \cos 2\Phi_N(t)\ ,$$

$$h_\times = -\frac{4(1+z)\mathcal{M}}{D_L}\left[\pi(1+z)\mathcal{M}f\right]^{2/3}\cos\iota\sin 2\Phi_N(t)\ ,$$

$$\Phi_N(t) = \Phi_c - \left[\frac{t_c-t}{5(1+z)\mathcal{M}}\right]^{5/8}\ , \qquad f \equiv \frac{1}{\pi}\frac{d\Phi_N}{dt}\ . \quad (1)$$

Here Φ_N is the orbital phase, f is the GW frequency, and $\mathcal{M} = m_1^{3/5}m_2^{3/5}/(m_1+m_2)^{1/5}$ is the binary's "chirp mass," which sets the rate at which the frequency changes. We use units with $G = 1 = c$; handy conversion factors are $M_\odot \equiv GM_\odot/c^2 = 1.47\,\mathrm{km}$, and $M_\odot \equiv GM_\odot/c^3 = 4.92\times10^{-6}\,\mathrm{seconds}$. The angle ι describes the inclination of the binary's orbital plane to our line-of-sight: $\cos\iota = \hat{\mathbf{L}}\cdot\hat{\mathbf{n}}$, where $\hat{\mathbf{L}}$ is the unit vector normal to the binary's orbital plane, and $\hat{\mathbf{n}}$ is the unit vector along the line-of-sight to the binary. The parameters t_c and Φ_c are the time and orbital phase when the frequency f diverges in this model. (In reality, we expect finite size effects to substantially impact the waveform before this.)

A given detector measures a linear combination of these polarizations:

$$h_{\rm meas} = F_+(\theta,\phi,\psi)h_+ + F_\times(\theta,\phi,\psi)h_\times\ , \quad (2)$$

where θ and ϕ describe the binary's position on the sky, and the "polarization angle" ψ sets the inclination of the components of $\hat{\mathbf{L}}$ orthogonal to $\hat{\mathbf{n}}$. The angles ι and ψ fully specify the orientation vector $\hat{\mathbf{L}}$. For a particular detector geometry, the antenna functions F_+ and $F_\times$ can be found in Thorne (1987). (In Sec. 2.2 we give a general form for the gravitational waveform, without appealing to a specific detector, following the analysis of Cutler & Flanagan 1994, hereafter abbreviated CF94.)

Several features of Eqs. (1) and (2) are worth commenting upon. First, note that the phase depends on the *redshifted* chirp mass. Measuring phase thus determines the redshifted chirp mass (Finn & Chernoff 1993). To understand why we measure redshifted chirp mass, note that $\mathcal{M}$ controls how fast the frequency evolves: using Eq. (1), we find $\dot{f} \propto f^{11/3}\mathcal{M}^{5/3}$. One can thus regard the chirp mass as entering the system's dynamics as a chirp time $\tau_c = G\mathcal{M}/c^3$. For a source at cosmological distance, this timescale is redshifted; the chirp mass we infer is likewise redshifted. Redshift and chirp mass are inextricably degenerate. This remains true even when higher order effects (see, e.g., Blanchet 2006) are taken into account: parameters describing a binary impact its evolutionary dynamics as timescales which suffer cosmological redshift; the inferred values of those parameters are thus redshifted. *GW observations on their own cannot directly determine a source's redshift.*

We next note that the amplitude depends on $(1+z)\mathcal{M}$, the angles (θ,ϕ,ι,ψ), and the luminosity distance D_L. Measuring a GW amplitude thus measures some tangled combination of these parameters. By measuring the phase, we measure the redshifted chirp mass sufficiently well that $(1+z)\mathcal{M}$ essentially decouples from the amplitude. More concretely, matched filtering the datastream with waveform templates should allow us to determine the phase with fractional accuracy $\delta\Phi/\Phi \sim 1/[(\text{signal-to-noise})\times(\text{number of measured cycles})]$; $(1+z)\mathcal{M}$ should be measured with similar fractional accuracy. Since NS-NS binaries will radiate roughly 10^4 cycles in band, and NS-BH binaries roughly 10^3 cycles, the accuracy with which phase and redshifted chirp mass can be determined should be exquisite.

Although $(1+z)\mathcal{M}$ therefore decouples from the amplitude, the distance, position, and orientation angles remain highly coupled. If our goal is to determine the source distance, we must break the degeneracy that the amplitude's functional form sets among these parameters. For sources that we measure with ground-based GW detectors, one way to break these degeneracies is to measure the waves with multiple detectors. Studies (Sylvestre 2004; Cavalier et al. 2006; Blair et al. 2008)

[6] http://www.gravity.uwa.edu.au
[7] http://gw.icrr.u-tokyo.ac.jp:8888/lcgt/

have shown that doing so allows us to determine the position of a merging binary to within a few degrees, giving us some information about the source's distance and inclination.

Perhaps the best way to break these degeneracies would be to measure the event electromagnetically. An EM signature is almost certain to pin down the event's position far more accurately than GWs alone. The position angles then effectively decouple, much as the redshifted chirp mass decoupled. Using multiple detectors, we can then determine the source's orientation and its distance. This gives us a direct, calibration-free measure of the distance to a cosmic event. The EM signature may also provide us with the event's redshift, breaking the GW degeneracy of redshift and intrinsic parameters, and *directly measuring both an event's distance and redshift*. In addition, if modeling or observations give us evidence for beaming of the SHB emission, that could strongly constrain the source inclination.

1.3. *This work and previous analysis*

Our goal is to assess how well we can determine the distance D_L to SHBs under the assumption that they are associated with inspiral GWs. We consider both NS-NS and NS-BH mergers as generators of SHBs, and consider several plausible advanced detector networks: the current LIGO/Virgo network, upgraded to advanced sensitivity; LIGO/Virgo plus the proposed Australian AIGO; LIGO/Virgo plus the proposed Japanese LCGT; and LIGO/Virgo plus AIGO plus LCGT.

The engine of our analysis is the construction of a probability function describing how measured source parameters (focusing in particular on D_L) should be distributed following GW measurement. Briefly, this formalism builds a probability distribution for the parameters $\boldsymbol{\theta}$ of a source's waveform. (Components θ^a of the vector $\boldsymbol{\theta}$ are physical parameters such as a binary's masses, distance, sky position angles, etc.) Consider one detector which measures a datastream $s(t)$, containing noise $n(t)$ and a GW signal $h(t,\hat{\boldsymbol{\theta}})$, where $\hat{\boldsymbol{\theta}}$ describes the source's "true" parameters. In the language of Finn (1992), we assume "detection" has already occurred; our goal in this paper is to focus on the complementary problem of "measurement." Our aim is to assess how well we can determine these parameters from our data.

As shown by Finn (1992), given a model for our signal $h(t,\boldsymbol{\theta})$, and assuming that the noise statistics are Gaussian, the probability that the parameters $\boldsymbol{\theta}$ describes the data s is

$$p(\boldsymbol{\theta}|s) = p_0(\boldsymbol{\theta}) \exp\left[-\left((h(\boldsymbol{\theta})-s)|(h(\boldsymbol{\theta})-s)\right)/2\right] . \quad (3)$$

The inner product $(a|b)$ describes the noise weighted cross-correlation of $a(t)$ with $b(t)$, and is defined precisely below. The distribution $p_0(\boldsymbol{\theta})$ is a *prior probability distribution*; it encapsulates what we know about our signal prior to measurement. We define $\tilde{\boldsymbol{\theta}}$ to be the parameters that maximize Eq. (3).

DHHJ06 did a first pass on the analysis which we describe here. They used an expansion in the variables $(\boldsymbol{\theta}-\hat{\boldsymbol{\theta}})$ of the exponential to second order in the limit of large signal-to-noise ratio (SNR), which we will henceforth refer to as the "Gaussian" approximation (cf. Finn 1992):

$$\exp\left[-\left(h(\boldsymbol{\theta})-s|h(\boldsymbol{\theta})-s\right)/2\right] \simeq \exp\left[-\frac{1}{2}\left(\frac{\partial h}{\partial\theta^a}\middle|\frac{\partial h}{\partial\theta^b}\right)\delta\theta^a\delta\theta^b\right] , \quad (4)$$

where $\delta\theta^a = \theta^a - \hat{\theta}^a$. In this limit, $\tilde{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}$ (at least for uniform priors). The matrix

$$\Gamma_{ab} \equiv \left(\frac{\partial h}{\partial\theta^a}\middle|\frac{\partial h}{\partial\theta^b}\right) \quad (5)$$

is the *Fisher information matrix*; its inverse Σ^{ab} is the covariance matrix. Diagonal entries Σ^{aa} describe the variance of parameter θ^a; off-diagonal entries describe correlations between different parameters.

The Gaussian approximation to Eq. (3) is known to be accurate when the SNR is large; however, it is not clear what "large" really means. Given current binary coalescence rate estimates, it is expected that most events will come from $D_L \sim$ a few $\times\, 100\,$Mpc. In such cases, we can expect an advanced detector SNR ~ 10. It is unclear that this value is sufficiently high such that the "large SNR" approximation is appropriate. Indeed, some of the results in DHHJ06 seem somewhat anomalous. For example, it had been expected that an Australian detector would have a disproportionate impact on the network's ability to determine distance and inclination for many events, since a southern hemisphere detector's antenna functions would be substantially different than its northern cousins, offering excellent constraints on an event's position and inclination. In fact, DHHJ06 found that a southern detector's main impact was simply to add additional SNR to a measurement—helpful, but not disproportionately so. We are concerned that this might be an artifact of the Gaussian approximation used in DHHJ06.

In this analysis we do *not* use this Gaussian approximation, but instead use Markov Chain Monte Carlo (MCMC) techniques (in particular, the Metropolis-Hastings algorithm) to explore our parameter distributions. The details of this technique are summarized in Sec. 3, and described in detail in Lewis & Bridle (2002). We find that MCMC techniques indeed indicate that the Gaussian approximation to Eq. (3) is failing in its estimate of a system's extrinsic parameters (though it appears to do quite well for intrinisic parameters such as masses). In particular, we find that an Australian detector would, in fact, have an outstanding impact on our ability to use SHB GWs as standard sirens. In addition, DHHJ06 restrict themselves to NS-NS binaries, observed with a four detector network. In what follows we relax both of these assumptions.

1.4. *Organization of this paper*

We begin in Sec. 2 by summarizing how GWs encode the distance to a coalescing binary. We first describe the post-Newtonian (PN) gravitational waveform we use in Sec. 2.1, and then describe how that wave interacts with and can be measured by a network of detectors in Sec. 2.2. Our discussion of the network-wave interaction is heavily based on the notation and formalism used in Sec. 4 of CF94, as well as the analysis of Anderson et al. (2001). Section 2.2 is sufficiently dense that we summarize its major points in Sec. 2.3 before concluding, in Sec. 2.4, with a description of the various GW detectors which we include in our analysis.

We outline our parameter estimation formalism in Sec. 3. In Sec. 3.1 we describe in more detail how to construct the probability distributions describing parameter measurement mentioned above (Sec. 1.3). We then give, in Sec. 3.2, a brief description of our selection procedure based on SNR detection thresholds. This procedure sets physically motivated priors for some of our parameters. The Markov Chain Monte Carlo technique we use to explore this function is described in Sec. 3.3. How one appropriately averages this distribution to give "noise averaged" results (and to compare with previous literature) is discussed in Sec. 3.4.

In Sec. 4 we discuss the validation of our code. We begin by attempting to reproduce some of the key results on distance measurement presented in CF94. Because of the rather different techniques used by Cutler & Flanagan, we do not expect exact agreement. It is reassuring to find, nonetheless, that we can reconstruct with very good accuracy all of the major features of their analysis. We then examine how these results change as we vary the amplitude (moving a fiducial test binary to smaller and larger distances), as we vary the number of detectors in our network, and as we vary the source's inclination.

Our main results are given in Sec. 5. We consider several different plausible detector networks and examine measurement errors for two "fiducial" binary systems, comprising either two neutron stars (NS-NS) with physical masses (i.e., rest frame masses, not including redshift effects) of $m_1 = m_2 = 1.4\,M_\odot$, or a neutron star and black hole (NS-BH) system with physical masses $m_1 = 1.4\,M_\odot$ and $m_2 = 10\,M_\odot$. Assuming a constant comoving cosmological density, we distribute potential GW-SHB events on the sky, and select from this distribution using a detection threshold criteria set for the entire GW detector network. We summarize some implications of our results in Sec. 6. A more in-depth discussion of these implications, particularly with regard to what they imply for the ability of standard sirens to measure the expansion of the universe, will be presented in a companion paper.

Throughout this paper we use units with $G = 1$, $c = 1$. Because we often refer to "redshifted" masses, we define the shorthand $m_z = (1+z)m$ for any mass parameter m.

2. MEASURING GRAVITATIONAL WAVES FROM INSPIRALING BINARIES

In this section we review the GW description used in our analysis, the formalism describing how these waves interact with a network of detectors, and the properties of the detectors.

2.1. *GWs from inspiraling binaries*

The inspiral and merger of a compact binary's members can be divided into three consecutive phases. The first and longest is a gradual adiabatic *inspiral*, when the members slowly spiral towards one another driven by the loss of orbital energy and angular momentum due to radiative backreaction. Post-Newtonian (PN) techniques (an expansion in the gravitational potential M/r, or equivalently for bound systems, the orbital speed v^2) allow a binary's evolution, and its emitted GWs, to be modeled analytically to high order; see Blanchet (2006) for a review. When the bodies come close together, the PN expansion is no longer valid, and direct numerical calculation is required. Recent breakthroughs in numerical relativity are making it possible to fully model the strong-field, dynamical *merger* of the two bodies into one; see Pretorius (2005), Shibata & Uryū (2006), and Etienne et al. (2008) for discussion. If the end state is a single black hole, the final waves from the system should be described by a *ringdown* as the black hole settles down to the Kerr spacetime solution. Note that much success has been achieved in analytically modeling the *entire* coalescence process using the "effective one body" technique. Initially proposed in Buonanno & Damour (1999), recent results show good agreement over the entire coalescence with the most accurately available numerical simulations (see Buonanno et al. 2009).

In this work we are concerned only with the inspiral phase, and will accordingly use the PN waveform to describe our waves. In particular, we use the so-called "restricted" PN waveform. Following CF94, the waveform of a binary inspiral may be written schematically

$$h(t) = \mathrm{Re}\left(\sum_{x,m} h^x_m(t) e^{im\Phi_{\rm orb}(t)}\right) . \qquad (6)$$

Here x indicates the PN order [h^x is computed to $O(v^{2x})$ in orbital speed], m denotes harmonic order (e.g., $m = 2$ is quadrupole), and $\Phi_{\rm orb}(t) = \int^t \Omega(t')dt'$ is orbital phase [with $\Omega(t)$ the orbital angular frequency]. The "restricted" waveform neglects all PN amplitude terms beyond the leading one, and considers only the dominant $m = 2$ contribution to the phase. The phase is itself computed to high PN order. As argued in CF94, these restrictions are useful because accumulated phase information has a greater impact on measurement accuracies than information in the amplitude.

Let the unit vector $\hat{\mathbf{n}}$ point to a binary on the sky (so that the waves propagate to us along $-\hat{\mathbf{n}}$), and let the unit vector $\hat{\mathbf{L}}$ denote the normal along the binary's orbital angular momentum. The waveform is fully described by the two polarizations:

$$\begin{aligned} h_+(t) &= \frac{2\mathcal{M}_z}{D_L}\left[\pi\mathcal{M}_z f(t)\right]^{2/3}\left[1+(\hat{\mathbf{L}}\cdot\hat{\mathbf{n}})^2\right]\cos[\Phi(t)] , \\ &\equiv \frac{4\mathcal{M}_z}{D_L}\left[\pi\mathcal{M}_z f(t)\right]^{2/3}\mathcal{A}_+(\hat{\mathbf{n}},\hat{\mathbf{L}})\cos[\Phi(t)] ; \qquad (7) \\ h_\times(t) &= -\frac{4\mathcal{M}_z}{D_L}[\pi\mathcal{M}_z f(t)]^{2/3}(\hat{\mathbf{L}}\cdot\hat{\mathbf{n}})\sin[\Phi(t)] , \\ &\equiv \frac{4\mathcal{M}_z}{D_L}\left[\pi\mathcal{M}_z f(t)\right]^{2/3}\mathcal{A}_\times(\hat{\mathbf{n}},\hat{\mathbf{L}})\sin[\Phi(t)] . \qquad (8) \end{aligned}$$

Equations (7) and (8) are nearly identical to those given in Eq. (1). In particular, $\mathcal{M}_z$ is the binary's redshifted chirp mass, D_L is the luminosity distance to the binary, and we have explicitly defined the inclination angle $\cos\iota$ in terms of the vectors $\hat{\mathbf{n}}$ and $\hat{\mathbf{L}}$. For later convenience, we have moved all dependence on sky position and orientation into the functions $\mathcal{A}_{+,\times}$. In Sec. 2.2, we discuss how these polarizations enter the GW field more generally, and how they interact with our detectors.

A major difference in these forms of h_+ and $h_\times$, as compared to Eq. (1), is that we now compute the GW phase to higher order:

$$\Phi(t) = 2\pi\int f(t')\,dt' = 2\pi\int \frac{f}{df/dt}df , \qquad (9)$$

using the 2nd-post-Newtonian (2PN) frequency sweep (Blanchet et al. 1995)

$$\frac{df}{dt} = \frac{96}{5}\pi^{8/3}\mathcal{M}_z^{5/3}f^{11/3}\left[1 - \left(\frac{743}{336}+\frac{11}{4}\eta\right)(\pi M_z f)^{2/3} + (4\pi)(\pi M_z f) + \left(\frac{34103}{18144}+\frac{13661}{2016}\eta+\frac{59}{18}\eta^2\right)(\pi M_z f)^{4/3}\right] . \quad (10)$$

Higher order results for df/dt are now known (Blanchet et al. 2002a,b, 2004), but 2PN order will be adequate for our purposes. Since distance measurements depend on accurate amplitude determination, we do not need a highly refined model of the wave's phase. The rate of sweep is dominantly determined by the chirp mass, but there is an important correction due to $\eta = \mu/M = m_1 m_2/(m_1+m_2)^2$, the reduced mass ratio. Note that η is not redshifted; both μ and M (the reduced mass and total mass, respectively) acquire $(1+z)$ corrections, so their ratio is the same at all z. Accurate measurement of the frequency sweep can thus determine both $\mathcal{M}_z$ and η (or $\mathcal{M}_z$ and μ_z), potentially determining the redshifted masses of each member of the binary.

We will often find it useful to work in the frequency domain, using the Fourier transform $\tilde{h}(f)$ rather than $h(t)$:

$$\tilde{h}(f) \equiv \int_{-\infty}^{\infty} e^{2\pi i f t} h(t)\, dt . \quad (11)$$

An approximate analytical result for $\tilde{h}(f)$ can be found using the stationary phase approximation (Finn & Chernoff 1993), which describes the Fourier transform when f changes slowly:

$$\tilde{h}_+(f) = \sqrt{\frac{5}{96}}\frac{\pi^{-2/3}\mathcal{M}_z^{5/6}}{D_L}\mathcal{A}_+ f^{-7/6}e^{i\Psi(f)} , \quad (12)$$

$$\tilde{h}_\times(f) = \sqrt{\frac{5}{96}}\frac{\pi^{-2/3}\mathcal{M}_z^{5/6}}{D_L}\mathcal{A}_\times f^{-7/6}e^{i\Psi(f)-i\pi/2} . \quad (13)$$

"Slowly" means that f does not change very much over a single wave period $1/f$, so that $(df/dt)/f \ll f$. The validity of this approximation for the waveforms we consider, at least until the last moments before merger, has been demonstrated in previous work (Droz et al. 1999). The phase function $\Psi(f)$ in Eqs. (12) and (13) is given by

$$\Psi(f) = 2\pi f t_c - \Phi_c - \frac{\pi}{4} + \frac{3}{128}(\pi\mathcal{M}f)^{-5/3}\times \left[1 + \frac{20}{9}\left(\frac{743}{336}+\frac{11}{4}\eta\right)(\pi M_z f)^{2/3} - 16\pi(\pi M_z f) + 10\left(\frac{3058673}{1016064}+\frac{5429}{1008}\eta+\frac{617}{144}\eta^2\right)(\pi M_z f)^{4/3}\right] . \quad (14)$$

As in Eq. (1), t_c is called the "time of coalescence" and defines the time at which f diverges within the PN framework; Φ_c is similarly the "phase at coalescence." Following the method employed in earlier works, we assume an abrupt (and unphysical) transition between the inspiral and merger phases at the so-called innermost stable circular orbit (ISCO), $f_{\rm ISCO} = (6\sqrt{6}\pi M_z)^{-1}$. For NS-NS, $f_{\rm ISCO}$ occurs at high frequencies where ground-based detectors have poor sensitivity; as such, we are confident that this abrupt transition has little impact on our results. For NS-BH, $f_{\rm ISCO}$ is likely to be in a band with good sensitivity. Better modeling of this transition could be important in this case.

In this analysis, we neglect effects which depend on spin. In general relativity, spin drives precession effects which can "color" the waveform in important ways, and which can in principle have important observational effects (see, e.g., Lang & Hughes 2006, van der Sluys et al. 2008). These effects are important when the dimensionless Kerr spin parameter is fairly large. Neutron stars are unlikely to spin fast enough for their angular momentum to drive interesting precessions during the time that they are in the band of GW detectors. To show this, write the moment of inertia of a neutron star as

$$I_{\rm NS} = \frac{2}{5}\kappa M_{\rm NS}R_{\rm NS}^2 , \quad (15)$$

where $M_{\rm NS}$ and $R_{\rm NS}$ are the star's mass and radius, and the parameter κ describes the extent to which its mass is centrally condensed (compared to a uniform sphere). Detailed calculations with different equations of state indicate $\kappa \sim 0.7$–1 [cf. Cook et al. (1994); this range comes the slowly rotating configurations in their Tables 12, 15, 18, and 21]. For a neutron star whose spin period is $P_{\rm NS}$, one finds that the Kerr parameter is given by

$$a_{\rm NS} = \frac{2\pi c I_{\rm NS}}{G M_{\rm NS}^2 P_{\rm NS}} \simeq 0.06\kappa\left(\frac{R_{\rm NS}}{12\,{\rm km}}\right)^2\left(\frac{1.4\,M_\odot}{M_{\rm NS}}\right)\left(\frac{10\,{\rm msec}}{P_{\rm NS}}\right) . \quad (16)$$

As long as the neutron star spin period is longer than ~ 10 msec, $a_{\rm NS}$ is small enough that spin effects can probably be safely neglected in our analysis. Strictly speaking, we *should* include spin in our models of BH-NS binaries; we leave this to a later analysis. We note that van der Sluys et al. (2008) included black hole spin effects in an analysis which did not assume known source position. They found that spin-induced modulations greatly improved the ability of GW detectors to localize a source, even when measured by only one detector. This suggests that, if position is known, spin modulations would likewise greatly improve our ability to measure source inclination and distance.

Our GWs depend on nine parameters: two mass parameters $\mathcal{M}_z$ and μ_z, two sky position angles (which set $\hat{\mathbf{n}}$), two orientation angles (which set $\hat{\mathbf{L}}$), time at coalescence t_c, phase at coalescence Φ_c, and luminosity distance to the source D_L. When sky position is known (e.g., from an SHB observation, as assumed in this analysis), the parameter set is reduced to seven: $\{\mathcal{M}_z, \mu_z, D_L, t_c, \cos\iota, \psi, \Phi_c\}$.

2.2. *Measurement of GWs by a detector network*

We now examine how the waves described in Sec. 2.1 interact with a network of detectors. We begin by introducing some conventions to describe our binary and our detectors. We do this geometrically, using vectors which describe our source's position and orientation, and our detectors' positions and orientations.

As described in the previous section, a source's sky position is given by a unit vector $\hat{\mathbf{n}}$ (which points from the center of the Earth to the binary), and its orientation is given by a unit vector $\hat{\mathbf{L}}$ (which points along the binary's orbital angular momentum). We construct from these a pair of axes which describe the binary's orbital plane:

$$\hat{\mathbf{X}} = \frac{\hat{\mathbf{n}} \times \hat{\mathbf{L}}}{|\hat{\mathbf{n}} \times \hat{\mathbf{L}}|} \ , \quad \hat{\mathbf{Y}} = -\frac{\hat{\mathbf{n}} \times \hat{\mathbf{X}}}{|\hat{\mathbf{n}} \times \hat{\mathbf{X}}|} \ . \tag{17}$$

These axes in turn are used to define the *polarization basis tensors*

$$\mathbf{e}^{+} = \hat{\mathbf{X}} \otimes \hat{\mathbf{X}} - \hat{\mathbf{Y}} \otimes \hat{\mathbf{Y}} \ , \tag{18}$$

$$\mathbf{e}^{\times} = \hat{\mathbf{X}} \otimes \hat{\mathbf{Y}} + \hat{\mathbf{Y}} \otimes \hat{\mathbf{X}} \ . \tag{19}$$

From these, the transverse-traceless metric perturbation which describes the GWs emitted by our source can be written

$$h_{ij} = h_{+} e^{+}_{ij} + h_{\times} e^{\times}_{ij} \ . \tag{20}$$

We next characterize the GW detectors. Each detector is an L-shaped interferometer whose arms define two-thirds of an orthonormal triple. Denote by $\hat{\mathbf{x}}_a$ and $\hat{\mathbf{y}}_a$ the unit vectors along the arms of the a-th detector in our network; we call these the x- and y-arms. (The vector $\hat{\mathbf{z}}_a = \hat{\mathbf{x}}_a \times \hat{\mathbf{y}}_a$ points radially from the center of the Earth to the detector's vertex.) These vectors define the *response tensor* for detector a:

$$D^{ij}_a = \frac{1}{2} \left[(\hat{\mathbf{x}}_a)^i (\hat{\mathbf{x}}_a)^j - (\hat{\mathbf{y}}_a)^i (\hat{\mathbf{y}}_a)^j \right] \ . \tag{21}$$

The response of detector a to a GW is given by

$$\begin{aligned} h_a &= D^{ij}_a h_{ij} \\ &\equiv e^{-2\pi i (\mathbf{n} \cdot \mathbf{x}_a) f} (F_{a,+} h_{+} + F_{a,\times} h_{\times}) \ . \end{aligned} \tag{22}$$

where $\mathbf{x_a}$ is the position of the ath detector and the factor $(\mathbf{n} \cdot \mathbf{x}_a)$ measures the time of flight between detector a and the coordinate origin. The second form of Eq. (22) shows how the antenna response functions introduced in Eq. (2) are built from the wave tensor and the response tensor.

Thus far this discussion has been completely frame-independent, in that we have defined all of our vectors and tensors without reference to a particular coordinate system. To facilitate an analysis of how inspiral GWs interact with a network of detectors, we now introduce a coordinate system for our detectors following the analysis of Anderson et al. (2001) [who in turn use the WGS-84 Earth model (Althouse et al. 1998)]. In this system, the Earth is taken to be an oblate ellipsoid with semi-major axis $a = 6.378137 \times 10^6$ meters, and semi-minor axis $b = 6.356752314 \times 10^6$ meters. Our coordinates are fixed relative to the center of the Earth. The x-axis (which points along $\mathbf{i}$) pierces the Earth at latitude 0° North, longitude 0° East (normal to the equator at the prime meridian); the y-axis (along $\mathbf{j}$) pierces the Earth at 0° North, 90° East (normal to the equator in the Indian ocean somewhat west of Indonesia); and the z-axis (along $\mathbf{k}$) pierces the Earth at 90° North (the North geographic pole).

A GW source at (θ, ϕ) on the celestial sphere has sky position vector $\hat{\mathbf{n}}$:

$$\hat{\mathbf{n}} = \sin\theta \cos\phi \mathbf{i} + \sin\theta \sin\phi \mathbf{j} + \cos\theta \mathbf{k} \ . \tag{23}$$

The *polarization angle*, ψ, is the angle (measured clockwise about $\hat{\mathbf{n}}$) from the orbit's line of nodes to the source's $\hat{\mathbf{X}}$-axis. In terms of these angles, the vectors $\hat{\mathbf{X}}$ and $\hat{\mathbf{Y}}$ are given by (Anderson et al. 2001)

$$\begin{aligned} \hat{\mathbf{X}} = &(\sin\phi\cos\psi - \sin\psi\cos\phi\cos\theta)\mathbf{i} \\ &-(\cos\phi\cos\psi + \sin\psi\sin\phi\cos\theta)\mathbf{j} + \sin\psi\sin\theta\mathbf{k} \ , \end{aligned} \tag{24}$$

$$\begin{aligned} \hat{\mathbf{Y}} = &(-\sin\phi\sin\psi - \cos\psi\cos\phi\cos\theta)\mathbf{i} \\ &+(\cos\phi\sin\psi - \cos\psi\sin\phi\cos\theta)\mathbf{j} + \cos\psi\sin\theta\mathbf{k} \ . \end{aligned} \tag{25}$$

Note that the angle ϕ is related to the right ascension, α, by $\alpha = \phi + \mathrm{GMST}$ (where GMST is the Greenwich mean sidereal time at which the signal arrives), and θ is related to the declination, δ, by $\delta = \pi/2 - \theta$ (cf. Anderson et al. 2001, Appendix B). Combining Eqs. (24) and (25) with Eqs. (18)–(20) allows us to write h_{ij} for a source in coordinates adapted to this problem.

We now similarly describe our detectors in terms of convenient coordinates. Detector a is at East longitude λ_a and North latitude φ_a (not to be confused with sky position angle ϕ). The unit vectors pointing East, North, and Up for this detector are

$$\mathbf{e}^{\mathrm{E}}_a = -\sin\lambda_a \mathbf{i} + \cos\lambda_a \mathbf{j} \ , \tag{26}$$

$$\mathbf{e}^{\mathrm{N}}_a = -\sin\varphi_a \cos\lambda_a \mathbf{i} - \sin\varphi_a \sin\lambda_a \mathbf{j} + \cos\varphi_a \mathbf{k} \ , \tag{27}$$

$$\mathbf{e}^{\mathrm{U}}_a = \cos\varphi_a \cos\lambda_a \mathbf{i} + \cos\varphi_a \sin\lambda_a \mathbf{j} - \cos\varphi_a \mathbf{k} \ . \tag{28}$$

The x-arm of detector a is oriented at angle Υ_a North of East, while its y-arm is at angle $\Upsilon_a + \pi/2$. Thanks to the Earth's oblateness, the x- and y-arms are tilted at angles $\omega^{x,y}_a$ to the vertical. The unit vectors $\hat{\mathbf{x}}_a$, $\hat{\mathbf{y}}_a$ can thus be written

$$\hat{\mathbf{x}}_a = \cos\omega^x_a \cos\Upsilon_a \mathbf{e}^{\mathrm{E}}_a + \cos\omega^x_a \sin\Upsilon_a \mathbf{e}^{\mathrm{N}}_a + \sin\omega^x_a \mathbf{e}^{\mathrm{U}} \ , \tag{29}$$

$$\hat{\mathbf{y}}_a = -\cos\omega^y_a \sin\Upsilon_a \mathbf{e}^{\mathrm{E}}_a + \cos\omega^y_a \cos\Upsilon_a \mathbf{e}^{\mathrm{N}}_a + \sin\omega^y_a \mathbf{e}^{\mathrm{U}} \ . \tag{30}$$

Combining Eqs. (29) and (30) with Eq. (21) allows us to write the response tensor for each detector in our network.

2.3. *Summary of the preceding section*

Section 2.2 is sufficiently dense that a brief summary may help clarify its key features, particularly with respect to the quantities we hope to measure. From Eq. (22), we find that each detector in our network measures a particular weighted sum of the two GW polarizations h_{+} and $h_{\times}$. We can rewrite the waveform detector a measures as

$$h_a = \frac{4\mathcal{M}_z}{D_L} \mathcal{A}_p \left[\pi \mathcal{M}_z f(t)\right]^{2/3} \cos\left[\Phi(t) + \Phi_p\right] \ , \tag{31}$$

where we have introduced detector a's "polarization amplitude"

$$\mathcal{A}_p = \sqrt{\left(F_{a,+}\mathcal{A}_{+}\right)^2 + \left(F_{a,\times}\mathcal{A}_{\times}\right)^2} \ , \tag{32}$$

and its "polarization phase"

$$\tan\Phi_p = \frac{F_{a,\times}\mathcal{A}_{\times}}{F_{a,+}\mathcal{A}_{+}} \ . \tag{33}$$

Cutler (1998) introduced these quantities to describe the response of the space-based *LISA* detector to a source's waves. As previously discussed, the intrinsic GW phase, $\Phi(t)$, is a strong function of the redshifted chirp mass, $\mathcal{M}_z$, the reduced mass, μ_z, the time of coalescence, t_c, and the phase at coalescence, Φ_c. Measuring the phase determines these four quantities, typically with very good accuracy.

Consider for a moment measurements by a single detector. The polarization amplitude and phase depend on the binary's sky position, (θ, ϕ) or $\hat{\mathbf{n}}$, and orientation, (ψ, ι) or $\hat{\mathbf{L}}$. [They also depend on detector position, (λ_a, φ_a), orientation, Υ_a, and tilt, (ω_a^x, ω_a^y). Since these angles are known and fixed, we ignore them in this discussion.] If the angles $(\theta, \phi, \psi, \iota)$ are not known, then a single detector cannot separate them, nor can it separate out the distance, D_L.

Multiple detectors can, at least in principle, separately determine these parameters. Each detector measures its own amplitude and polarization phase, and by combining their outputs, we can fit to the four unknown angles and the distance, D_L. Various works have analyzed how well this can be done assuming that the position and orientation are completely unknown (Sylvestre 2004; Cavalier et al. 2006; Blair et al. 2008). Van der Sluys et al. (2008) performed such an analysis for measurements of NS-BH binaries, including the effect of orbital precession induced by the black hole. These precessions effectively make the angles ι and ψ time dependent, breaking the degeneracy among these angles and D_L in a completely different way.

In what follows, we assume that an electromagnetic identification pins down the angles (θ, ϕ), so that they do not need to be determined from the GW data. We then face the substantially less challenging problem of determining ψ, ι, and D_L from the various waveforms, h_a, that we measure at the detectors. We will also examine the impact of a constraint on the inclination, ι. Long bursts are believed to be strongly collimated, emitting into jets with opening angles of just a few degrees. Less is known about the collimation of SHBs, but it is plausible that their emission may be primarily along a preferred axis (presumably the progenitor binary's orbital angular momentum axis). We are particularly interested in assessing the extent to which constraining ι to be within some Gaussian distribution with opening angle $\delta\iota \sim 25^\circ$ of the "face-on" configuration ($|\hat{\mathbf{n}} \cdot \hat{\mathbf{L}}| = 1$) improves our ability to determine the source's distance.

2.4. *GW detectors used in our analysis*

Here we briefly summarize the properties of the GW detectors that we consider:
LIGO: The Laser Interferometer Gravitational-wave Observatory consists of two 4 kilometer interferometers located in Hanford, Washington (US) and Livingston, Louisiana (US). These instruments have achieved their "initial" sensitivity goals; an upgrade to "advanced" sensitivity is expected to be completed around 2015. We show the projected advanced sensitivity in Fig. 1. Broadly speaking, the noise spectrum is dominated by seismic vibrations due to ground and human activity at $f < 10\,\mathrm{Hz}$, thermal vibrations in the band $10\,\mathrm{Hz} < f < 50\,\mathrm{Hz}$, and photon shot noise for $f > 50\,\mathrm{Hz}$. At present, the Hanford site also houses a 2 kilometer interferometer; this will be extended to a full 4 kilometer baseline in the advanced era. As such, the LIGO network will consist of two instruments at Hanford and one in Livingston. Because the two Hanford instruments are likely to have somewhat correlated noise (particularly at low frequencies), we assume only one instrument in Hanford for most of our analysis.

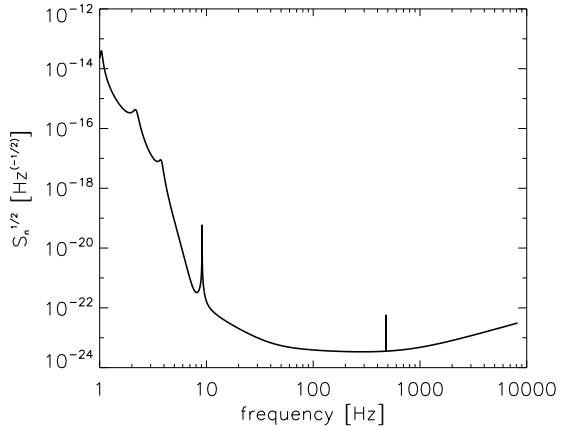

FIG. 1.— Projected noise curve for Advanced LIGO. We assume a lower frequency cut-off of 10 Hz.

Virgo: The Virgo detector near Pisa, Italy has slightly shorter arms than LIGO (3 kilometers), but should achieve similar advanced detector sensitivity, on roughly the same timescale as the LIGO detectors. For simplicity, we will assume that Virgo's sensitivity is the same as LIGO's in our analysis.

Our baseline detector network consists of the LIGO Hanford and Livingston sites, and Virgo; these are instruments which are running today, and will be upgraded over the next decade. We also examine the impact of adding two proposed interferometers to this network:
AIGO: The Australian International Gravitational Observatory is a proposed multikilometer interferometer that would be located in Gingin, Western Australia. AIGO's proposed site in Western Australia is particularly favourable due to low seismic and human activity.
LCGT: The Large-scale Cryogenic Gravitational-wave Telescope is a proposed multikilometer interferometer that would be located in the Kamioka observatory, 1 kilometer underground. This location takes advantage of the fact that local ground motions tend to decay rapidly as we move away from the Earth's surface. They also plan to use cryogenic cooling to reduce thermal noise.

As with Virgo, we will take the sensitivity of AIGO and LCGT to be the same as LIGO for our analysis. Table 1 gives the location and orientation of these detectors, needed to compute the response function for each member of our network.

3. ESTIMATION OF BINARY PARAMETERS

3.1. *Overview of formalism*

We now give a brief summary of parameter estimation. In particular, we lay the foundations for estimating measurement accuracies given a GW detector's noise spectrum. Further details are discussed in Finn (1992), Królak et al. (1993), and CF94.

The datastream of detector a, $s_a(t)$, has two major

TABLE 1
GW detectors we consider (positions and orientations)

Detector	East Long. λ	North Lat. φ	Orientation Υ	x-arm tilt ω^x	y-arm tilt ω^y
LIGO-Han	$-119.4°$	$46.5°$	$126°$	$(-6.20\times10^{-4})°$	$(1.25\times10^{-5})°$
LIGO-Liv	$-90.8°$	$30.6°$	$198°$	$(-3.12\times10^{-4})°$	$(-6.11\times10^{-4})°$
Virgo	$10.5°$	$43.6°$	$70°$	$0.0°$	$0.0°$
AIGO	$115.7°$	$-31.4°$	$0°$	$0.0°$	$0.0°$
LCGT	$137.3°$	$36.4°$	$25°$	$0.0°$	$0.0°$

contributions: The true GW signal $h_a(t;\hat{\boldsymbol{\theta}})$ (constructed by contracting a source's GW tensor h_{ij} with detector a's response tensor D_a^{ij}; cf. Sec. 2.2), and a specific realization of detector noise $n_a(t)$,

$$s_a(t) = h_a(t;\hat{\boldsymbol{\theta}}) + n_a(t)\ . \tag{34}$$

The incident gravitational wave strain is assumed to depend on (unknown) true parameters $\hat{\boldsymbol{\theta}}$. As in Sec. 1.3, $\hat{\boldsymbol{\theta}}$ is a vector whose components are binary parameters. Below we use a vector $\mathbf{s}$ whose components s_a are the datastreams of each detector (and likewise $\mathbf{h}$ and $\mathbf{n}$ are vectors whose components are the GW and noise content of each detector, respectively).

We assume the noise to be stationary, zero mean, and Gaussian. This lets us categorize it using the spectral density, as follows. First, define the noise correlation matrix:

$$\begin{aligned} C_n(\tau)_{ab} &= \langle n_a(t+\tau)n_b(t)\rangle - \langle n_a(t+\tau)\rangle\,\langle n_b(t)\rangle \\ &= \langle n_a(t+\tau)n_b(t)\rangle\ , \end{aligned} \tag{35}$$

where the angle brackets are ensemble averages over noise realizations, and the zero mean assumption gives us the simplified form on the second line. For $a = b$, this is the auto-correlation of detector a's noise; otherwise, it describes the correlation between detectors a and b. The (one-sided) power spectral density matrix is the Fourier transform of this:

$$S_n(f)_{ab} = 2\int_{-\infty}^{\infty} d\tau\, e^{2\pi i f\tau} C_n(\tau)_{ab}\ . \tag{36}$$

Note that this is defined for $f > 0$ only. For $a = b$, this is the spectral density of noise power in detector a; for $a \neq b$, it again describes correlations between detectors. From these definitions, one can show that

$$\langle \tilde{n}_a(f)\,\tilde{n}_b(f')^*\rangle = \frac{1}{2}\delta(f-f')S_n(f)_{ab}. \tag{37}$$

For Gaussian noise, this statistic completely characterizes our detector noise. No real detector is of course completely Gaussian, but by using multiple, widely-separated detectors non-Gaussian events can be rejected. In our analysis, we assume the detectors' noises are uncorrelated such that Eq. (37) becomes

$$\langle \tilde{n}_a(f)\,\tilde{n}_b(f')^*\rangle = \frac{1}{2}\delta_{ab}\delta(f-f')S_n(f)_a. \tag{38}$$

Finally, for simplicity we assume that $S_n(f)_a$ has the universal shape $S_n(f)$ projected for advanced LIGO. This spectrum is shown in Fig. 1.

The central quantity of interest in parameter estimation is the posterior probability distribution function (PDF) for $\boldsymbol{\theta}$ given detector output $\mathbf{s}$, which is defined as

$$p(\boldsymbol{\theta}\,|\,\mathbf{s}) = \mathcal{N}\,p^{(0)}(\boldsymbol{\theta})\mathcal{L}_{\rm TOT}(\mathbf{s}\,|\,\boldsymbol{\theta})\,. \tag{39}$$

(Note that we assume a GW has been detected.) In Eq. (39), $\mathcal{N}$ denotes a normalization constant, $p^{(0)}(\boldsymbol{\theta})$ is the PDF that represents the prior probability that a measured GW is described by the parameters $\boldsymbol{\theta}$, and $\mathcal{L}_{\rm TOT}(\mathbf{s}\,|\,\boldsymbol{\theta})$ is the total *likelihood function* (e.g., MacKay 2003). The likelihood function measures the relative conditional probability of observing a particular set of data $\mathbf{s}$ given that a measured signal $\mathbf{h}$ depending on some unknown set of parameters, $\boldsymbol{\theta}$ and noise, $\mathbf{n}$. Because we assume that the noise is independent and uncorrelated at each detector site, we may take the total likelihood function to be the product of the individual likelihood at each detector:

$$\mathcal{L}_{\rm TOT}(\mathbf{s}\,|\,\boldsymbol{\theta}) = \Pi_a\mathcal{L}_a(s_a\,|\,\boldsymbol{\theta})\ , \tag{40}$$

where $\mathcal{L}_a$, the likelihood computed at detector a, is given by (Finn 1992)

$$\mathcal{L}_a\,(s\,|\,\boldsymbol{\theta}) = \ e^{-\left(h_a(\boldsymbol{\theta})-s_a\,\left|\,h_a(\boldsymbol{\theta})-s_a\right.\right)/2}\,. \tag{41}$$

The inner product $(\ldots|\ldots)$ on the vector space of signals is defined as

$$(g|h) = 2\int_0^{\infty} df\,\frac{\tilde{g}^*(f)\tilde{h}(f)+\tilde{g}(f)\tilde{h}^*(f)}{S_n(f)}\,. \tag{42}$$

This definition means that the probability of the random noise $n(t)$ assuming a particular realization $n_0(t)$ is

$$p(n=n_0)\ \propto\ e^{-(n_0|n_0)/2}. \tag{43}$$

For clarity, we distinguish between various definitions of the signal-to-noise ratio (SNR). The *true* SNR at detector a, associated with a given instance of noise for a measurement at a particular detector, is defined as (see CF94)

$$\left(\frac{S}{N}\right)_{a,\rm true} = \frac{(h_a\,|\,s_a)}{\sqrt{(h_a\,|\,h_a)}}\,. \tag{44}$$

This is a random variable with Gaussian PDF of unit variance. For an ensemble of realizations of the detector noise n_a, the so-called *average* SNR at detector a is given by

$$\left(\frac{S}{N}\right)_{a,\rm ave} = \frac{(h_a|h_a)}{\text{rms }(h_a|n_a)} = (h_a|h_a)^{1/2}. \tag{45}$$

Consequently, we can define the combined *true* and *average* SNRs of a coherent network of detectors:

$$\left(\frac{S}{N}\right)_{\rm true} = \sqrt{\sum_a\left(\frac{S}{N}\right)^2_{a,\rm true}}\ , \tag{46}$$

and

$$\left(\frac{S}{N}\right)_{\rm ave} = \sqrt{\sum_a\left(\frac{S}{N}\right)^2_{a,\rm ave}}\ . \tag{47}$$

Estimating the parameter set $\boldsymbol{\theta}$ often uses the "maximum likelihood" method following either a Bayesian (Loredo 1989, Finn 1992, CF94, Poisson & Will 1995) or frequentist point of view (Królak et al. 1993, CF94). We do not attempt to review these differing philosophies, and instead refer to Appendix A2 of CF94 for detailed discussion. It is worth noting that, in the GW literature, the "maximum likelihood" or "maximum a posterior" are often interchangeably referred to as the "best-fit" parameters, which can be confusing. The maximum a posterior is the parameter set $\tilde{\boldsymbol{\theta}}_{\rm MAP}$ which maximizes the full posterior probability, Eq. (39); likewise, the maximum likelihood is the parameter set $\tilde{\boldsymbol{\theta}}_{\rm ML}$ which maximizes the likelihood function, Eq. (40).

Following the approach advocated by CF94, we introduce the Bayes estimator $\tilde{\theta}^i_{\rm BAYES}(\mathbf{s})$, defined as

$$\tilde{\theta}^i_{\rm BAYES}(\mathbf{s}) \equiv \int \theta^i \, p(\boldsymbol{\theta} \,|\, \mathbf{s}) d\boldsymbol{\theta} \, . \tag{48}$$

The integral is performed over the whole parameter set $\boldsymbol{\theta}$; $d\boldsymbol{\theta} = d\theta^1 d\theta^2 \ldots d\theta^n$. Similarly, we define the rms measurement errors $\Sigma^{ij}_{\rm BAYES}$

$$\Sigma^{ij}_{\rm BAYES} = \int (\theta^i - \tilde{\theta}^i_{\rm BAYES}) \, (\theta^j - \tilde{\theta}^j_{\rm BAYES}) \, p(\boldsymbol{\theta} \,|\, \mathbf{s}) d\boldsymbol{\theta}. \tag{49}$$

The meaning of the Bayes estimator can be understood as follows. Consider a single detector which records an arbitrarily large ensemble of GW signals. Within this ensemble, there will be a subset of signals with the same detector output $s(t)$. Although these outputs are identical, each actually corresponds to GW signals with different true parameters $\hat{\boldsymbol{\theta}}$; it just so happens that they also have different noise realizations $n(t)$ that conspire so that the sum $s(t)$ is the same. In this case, $\tilde{\theta}^i_{\rm BAYES}(\mathbf{s})$ represents the expectation of θ^i averaged over the ensemble of GW signals. The principle disadvantage of the Bayes estimator is computation cost, owing to the multi-dimensional integrals in Eqs. (48) and (49). In Sec. 3.3, we describe MCMC methods which efficiently approximate these integrals.

For large SNR it can be shown that the estimators $\tilde{\boldsymbol{\theta}}_{\rm ML}$, $\tilde{\boldsymbol{\theta}}_{\rm MAP}$, and $\tilde{\boldsymbol{\theta}}_{\rm BAYES}$ agree with one another (CF94). Furthermore, in this limit (and as mentioned in Sec. 1.3) Eq. (39) assumes a Gaussian form with respect to parameter errors, allowing us to estimate parameter variances as the inverse of a Fisher matrix. However, as illustrated in Sec. IVD of CF94, effects which scale nonlinearly with 1/SNR and prior information [represented by the distribution $p^{(0)}(\boldsymbol{\theta})$] contribute significantly at low SNR. The Gaussian approximation to Eq. (39) then tends to underestimate measurement errors by missing tails or multimodal structure in posterior distributions.

We emphasize that in this analysis we do not consider systematic errors that occur both due to limitations in our source model and due to gravitational lensing effects. A framework for analyzing systematic errors in GW measurements has recently been presented by Cutler & Vallisneri (2007). An important follow-on to this work will be to estimate systematic effects and determine whether they significantly change our conclusions.

3.2. *Binary Selection and Priors*

We now detail our method for generating a sample of detectable GW-SHB events. The selection is central to the use of GW-SHBs as standard sirens. In addition, it sets physically motivated priors for the parameters in our Bayesian framework. We employ a simple selection procedure based on detection thresholds. An implicit assumption of our method, detailed below, is that every SHB event is a GW candidate potentially detectable by a network of advanced GW interferometers.

We assume a constant comoving density (Peebles 1993, Hogg 1999) of GW-SHB events, in a ΛCDM Universe with $H_0 = 70.5$ km/sec/Mpc, $\Omega_\Lambda = 0.726$, and $\Omega_m = 0.2732$ (Komatsu et al. 2009); our results are insensitive to the precise details of the cosmology. We Monte-Carlo distribute 10^6 binaries uniformly in volume, with random sky positions and orientations, to redshift $z = 1$ (i.e., to roughly 6.6 Gpc). We compute the average SNR, Eq. (45), for each binary at each detector, and use Eq. (47) to compute the average total SNR for each network we consider. Since we assume prior knowledge of the merger time (by assumption that the inspiral is correlated with a SHB), we set a threshold SNR for the *total* detector network, $\mathrm{SNR}_{\rm total} = 7.5$ (see discussion in DHHJ06). This is somewhat reduced from the threshold we would set in the absence of a counterpart, since prior knowledge of merger time and source position reduces the number of search templates we need by a factor $\sim 10^5$ (Kochanek & Piran 1993, Owen 1996). By using the average SNR to set our threshold, we introduce a slight error into our analysis, since the true SNR will differ from the average. Some events which we identify as above threshold could be moved below threshold due to a measurement's particular noise realization. However, some sub-threshold events will likewise be moved above threshold due to the instance of noise. The net effect is not expected to be significant, and will be evaluated in future work.

Our choice of threshold selects detectable GW-SHB events for each detector network. To avoid confusion, we use the term "total detected binaries" to mean binaries which are detected by a network consisting of all five detectors—both LIGO sites, Virgo, AIGO, and LCGT. As one might imagine, this five detector network detects more binaries than one with four detectors, which in turn detects more than one of three detectors. Both the Southern Hemisphere detector AIGO and the Japanese detector LCGT substantially improve on the number of binaries detected, as compared to the two LIGO detectors plus Virgo. Assuming isotropy in SHB orientation (i.e., that all binary orientations are equally likely given an SHB), we find that a LIGO-LIGO-Virgo network detects 50% of the total detected binaries; LIGO-LIGO-Virgo-AIGO detects 74% of the total; and LIGO-LIGO-Virgo-LCGT detects 72% of the total. Figure 2 shows the location of detected binaries on the sky for LIGO and Virgo (upper left panel); LIGO, Virgo, and AIGO (upper right); LIGO, Virgo, and LCGT (lower left); and all detectors (lower right). Notice that networks which include LCGT tend to have rather uniform sky coverage, and that with AIGO, the quadrants with $\cos\theta > 0$, $\phi > \pi$ and with $\cos\theta < 0$, $\phi < \pi$ are covered particularly well.

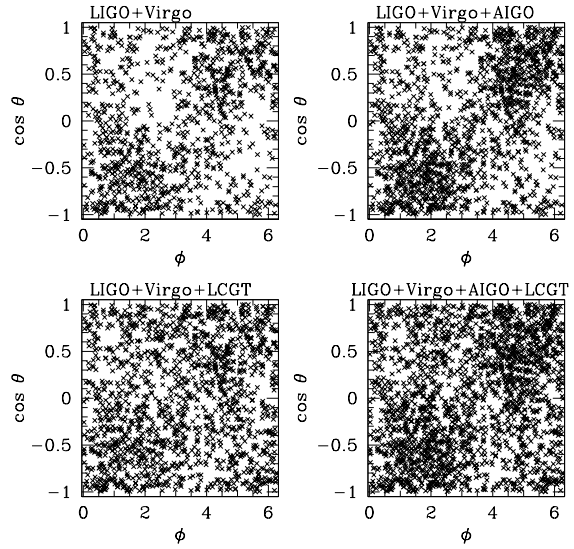

FIG. 2.— Detected NS-NS binaries for our various detector networks, as a function of sky position $(\cos\theta, \phi)$. The lower right panel shows all of the binaries detected by a five-detector network consisting of LIGO Hanford, LIGO Livingston, Virgo, AIGO, and LCGT; upper left is LIGO plus Virgo, with 50% of the five-detector events detected; upper right is LIGO, Virgo, and AIGO, with 74% of events detected; and lower left is LIGO, Virgo, and LCGT, with 72% of events detected. Notice that detected binaries are more uniformly distributed on the sky in networks that include LCGT, and that AIGO markedly improves the coverage in two of the sky's quadrants. Our coordinate ϕ is related to right ascension α by $\phi = \alpha -$GMST, where GMST is Greenwich Mean Sidereal Time; θ is related to declination δ by $\theta = \pi/2 - \delta$.

Our selection method implicitly sets a prior distribution on our parameters. For example, the thresholding procedure results in a significant bias in detected events toward face-on binaries, with $\hat{\mathbf{L}}\cdot\hat{\mathbf{n}} \rightarrow \pm 1$. Figure 3 shows the 2-D sample of detectable NS-NS binaries for the parameters $(\cos\iota, D_L)$. Since we use an unrealistic mass distribution ($1.4\,M_\odot$–$1.4\,M_\odot$ NS-NS and $1.4\,M_\odot$–$10\,M_\odot$ NS-BH binaries), instead of a more astrophysically realistic distribution, the implicit mass prior is uninteresting. In addition, we assume uniform prior densities for the parameter ψ between $(0, \pi)$. Figure 4 presents the average total SNR versus the true D_L of our sample of detectable NS-NS and NS-BH binaries for our "full" network (LIGO, Virgo, AIGO, LCGT). Very few detected binaries have SNR above 30 for NS-NS, and above 70 for NS-BH. It is interesting to note the different detectable ranges in luminosity distance between the two populations, with the NS-BH binaries detectable to over twice the distance.

We are also interested in seeing the impact that prior knowledge of SHB collimation may have on our ability to measure characteristics of these events. We note that, to date, there exist only two tentative observations which suggest that SHBs may have collimated emission (Grupe et al. 2006, Burrows et al. 2006, Soderberg et al. 2006); we therefore present results for both moderate collimation and for isotropic SHB emission. To obtain a sample of beamed SHBs, we assume that the burst emission is collimated along the orbital angular momentum axis, where baryon loading is minimized. Following DHHJ06, we use a distribution for $\cos\iota \equiv v$ of $dP/dv \propto \exp[-(1-v)^2/2\sigma_v^2]$, with $\sigma_v = 0.05$. This corresponds to a beamed binary population with one sigma (68%) of its distribution having an opening jet angle within roughly 25°. As Fig. 5 shows, we construct a beamed subsample by selecting events from the total sample of detected events such that its final distribution in inclination angle follows dP/dv. Future joint measurements of SHBs and GW-driven inspirals should enable us to constrain the beaming angles by comparing the measured rates of these two populations.

3.3. *Markov Chain Monte Carlo approach*

As mentioned in Sec. 3.1, the principle disadvantage of the Bayes estimators $\tilde{\theta}^i_{\rm BAYES}$ and $\Sigma^{ij}_{\rm BAYES}$ is the high computational cost of evaluating the multi-dimensional integrals which define them, Eqs. (48) and (49). To get around this difficulty, we use Markov Chain Monte Carlo (MCMC) methods to explore the PDFs describing the seven parameters $\{\mathcal{M}_z, \mu, D_L, \cos\iota, \psi, t_c, \Phi_c\}$. MCMC methods are widely used in diverse astrophysical applications, ranging from high precision cosmology (e.g. Dunkley et al. 2009, Sievers et al. 2009) to extrasolar planet studies (e.g. Ford 2005, Winn et al. 2007). They have seen an increase in use in GW measurement and parameter estimation studies in recent years, both for the case of analysis by the space-based detector *LISA* (e.g., Stroeer et al. 2006, Wickham et al. 2006, Cornish & Porter 2007, Porter & Cornish 2008), and for analysis by a network of ground-based detectors at initial sensitivity (e.g., Röver et al. 2007, van der Sluys et al. 2008). In this section, we briefly introduce MCMC methods, focusing in particular on the Metropolis-Hastings algorithm (Metropolis et al. 1953, Hastings 1970). Our goal is to describe the workings of our algorithm, rather than to review MCMC methods generally; we defer detailed proofs to the literature on this subject (e.g. Neal 1993, Gilks et al. 1996, MacKay 2003). Our discussion owes much to the helpful background provided by Sec. III of Christensen et al. (2004).

As the "Monte Carlo" part of its name implies, the central idea of MCMC is to generate a random sequence of parameter states that sample the posterior distribution, $p(\boldsymbol{\theta}|\mathbf{s})$. Let the nth sample in the sequence be $\boldsymbol{\theta}^{(n)}$. Then, if one draws a total of N random samples, Eqs.

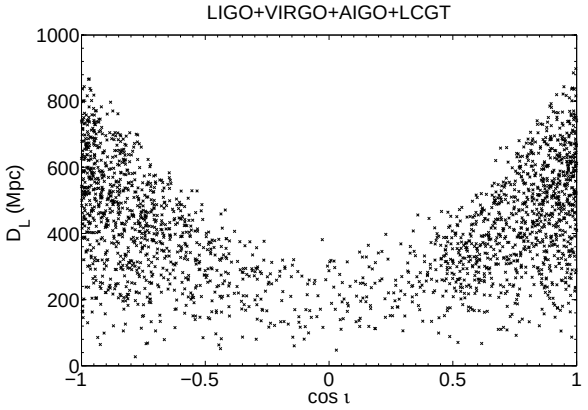

FIG. 3.— The 2-D marginalized prior distribution in luminosity distance, D_L, and cosine inclination, $\cos\iota$. Each point represents a detected NS-NS binary for a network comprising all five detectors. Notice the bias toward detecting face-on binaries ($\cos\iota \rightarrow \pm 1$); these are detected to much larger distances than edge-on ($\cos\iota \rightarrow 0$).

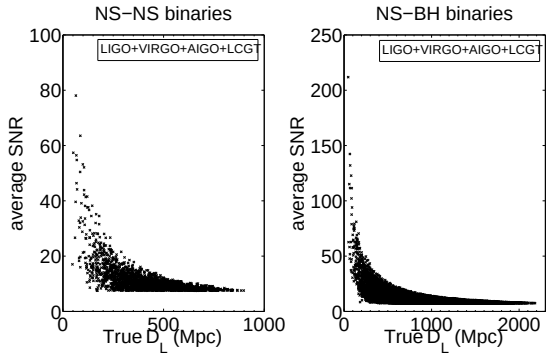

FIG. 4.— Average SNR versus luminosity distance of the total detected NS-NS and NS-BH binaries using a network of both LIGO detectors, Virgo, AIGO, and LCGT. The left panel shows the total detected NS-NS binaries (one point with SNR above 100 is omitted); the right panel shows the total detected NS-BH binaries (one point with SNR above 350 is omitted). Notice the different scales of both axes between the two panels. In particular, NS-BH binaries are detected to more than twice the distance.

(48) and (49) can be approximated as simple sample averages:

$$\tilde{\theta}^i_{\rm BAYES} \simeq \frac{1}{N}\sum_{n=1}^{N}(\theta^i)^{(n)} \ , \tag{50}$$

$$\Sigma^{ij}_{\rm BAYES} \simeq \frac{1}{N}\sum_{n=1}^{N}\left(\tilde{\theta}^i_{\rm BAYES} - (\theta^i)^{(n)}\right)\left(\tilde{\theta}^j_{\rm BAYES} - (\theta^j)^{(n)}\right) \ . \tag{51}$$

The key to making this technique work is drawing a sequence that represents the posterior PDF. The Metropolis-Hastings algorithm, which we now briefly describe, ensures that, at least asymptotically, the probability of selecting a sample $\boldsymbol{\theta}^{(n)}$ is determined by the posterior distribution $p(\boldsymbol{\theta}|\mathbf{s})$.

Suppose we are currently at step n in our chain of sequences; the zeroth step, $\boldsymbol{\theta}^{(0)}$, can be selected arbitrarily. We choose a candidate for the next sequence in our chain, $\boldsymbol{\theta}'$, from the *proposal distribution* $Q(\boldsymbol{\theta}';\boldsymbol{\theta}^{(n)})$. (Note that our choice for this function can have a large impact on how effectively the chain we build converges to a representation of the posterior PDF; we describe our particular choice later in this section.) We then accept (or reject) $\boldsymbol{\theta}'$ by computing the *acceptance probability*

$$a(\boldsymbol{\theta}';\boldsymbol{\theta}^{(n)}) = \frac{P(\boldsymbol{\theta}')Q(\boldsymbol{\theta}^{(n)};\boldsymbol{\theta}')}{P(\boldsymbol{\theta}^{(n)})Q(\boldsymbol{\theta}';\boldsymbol{\theta}^{(n)})} \ . \tag{52}$$

where $P(\boldsymbol{\theta})$ is shorthand for the posterior PDF $p(\boldsymbol{\theta}|\mathbf{s})$. We also generate a uniform random deviate $u \in [0,1]$. If $u \leq a(\boldsymbol{\theta}';\boldsymbol{\theta}^n)$, then we *accept* this step, putting $\boldsymbol{\theta}^{(n+1)} = \boldsymbol{\theta}'$. If $u > a(\boldsymbol{\theta}';\boldsymbol{\theta}^n)$, then this step is *rejected*, and we put $\boldsymbol{\theta}^{(n+1)} = \boldsymbol{\theta}^{(n)}$. Note that sample $\boldsymbol{\theta}^{(n+1)}$ only depends upon the previous sample, $\boldsymbol{\theta}^{(n)}$. This is the defining condition of a Markov sequence, and explains the "Markov Chain" part of MCMC. The literature on this subject demonstrates that the Markov chain so constructed explores the full range of the sample space spanned by $p(\boldsymbol{\theta}\,|\,\mathbf{s})$. We point the reader to Christensen et al. (2004) for a pedagogical review, and to Neal (1993) and Gilks et al. (1996) for much more in-depth discussion.

For completeness, we list several standard concerns of MCMC methods. First, as already mentioned in passing, this algorithm's performance depends on the choice of proposal density $Q(\boldsymbol{\theta}';\boldsymbol{\theta})$. The efficiency of the method increases when the proposal distribution resembles the underlying probability distributions, particularly when treating strongly degenerate parameters. Second, a certain so-called "burn-in" period is needed in order for a chain to equilibriate to its invariant distribution. This period allows the chain to explore all regions of high probability in the posterior PDF; it is especially important if the chain begins in some parameter region of low probability. Finally, a chain will of course necessarily be finite in any practical computation. Understanding the convergence properties of finite chains is central to ensuring that the samples we generate are sufficiently large, and are effectively independent. Some errors enter because of correlations between successive elements of the chain and the shot noise. We are especially concerned with establishing that the MCMC chain has fully sampled the distribution, and is not "stuck" on some parameter island of high probability.

The MCMC algorithm which we use is based on a generic version of CosmoMC[8], described in Lewis & Bridle (2002). We state here only the main features of our code; further details of our MCMC method are discussed at length in Lewis & Bridle (2002). As Eq. (52) illustrates, computing $a(\boldsymbol{\theta}';\boldsymbol{\theta}^{(n)})$ requires evaluating the posterior PDF $p(\boldsymbol{\theta}\,|\,\mathbf{s})$ for any sample $\boldsymbol{\theta}^{(n)}$. This in turn requires that we calculate both the prior distribution $p^{(0)}(\boldsymbol{\theta})$, and the likelihood function $\mathcal{L}_{\rm TOT}(\mathbf{s}\,|\,\boldsymbol{\theta})$ [cf. Eq. (40)] for $\boldsymbol{\theta} = \boldsymbol{\theta}^{(n)}$. We now briefly emphasize key aspects of our approach.

As already mentioned, the total likelihood $\mathcal{L}_{\rm TOT}(\mathbf{s}\,|\,\boldsymbol{\theta})$ depends upon the individual likelihoods at each detector, $\mathcal{L}_a(s_a\,|\,\boldsymbol{\theta})$. By Eq. (41), we see that each $\mathcal{L}_a(s_a\,|\,\boldsymbol{\theta})$ is a function of the detector's output $s_a(t)$, the predicted waveform $h_a(\boldsymbol{\theta}^{(n)})$ at $\boldsymbol{\theta}^{(n)}$, and the noise spectral density $S_n(f)$. For a given source (i.e., a given true parameterization $\hat{\boldsymbol{\theta}}$), we need only compute $s_a(t)$ once [cf. Eq. (34)]. When computing steps in our chain, we must then com-

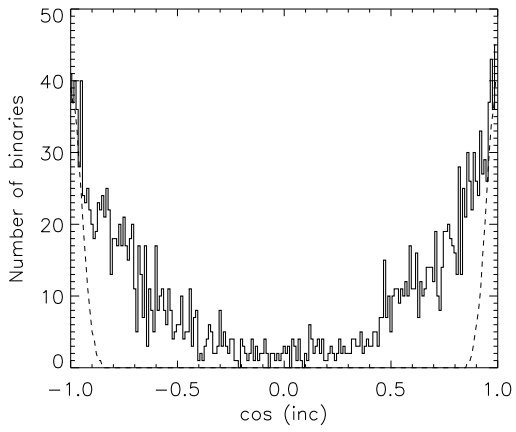

FIG. 5.— Distributions in inclination, $\cos\iota$, for the subsample of NS-NS binaries for which we assume SHB collimation (dashed line) and for the full sample of total detected binaries with isotropic distribution in inclination (solid line). The beamed subsample of binaries have a distribution in $\cos\iota \equiv v$ given by $dP/dv \propto \exp[-(1-v)^2/2\sigma_v^2]$, with $\sigma_v = 0.05$.

[8] See http://cosmologist.info/cosmomc/

pute the model signal $h_a(\boldsymbol{\theta}^{(n)})$ at each sample $\boldsymbol{\theta}^{(n)}$.

Priors play a crucial role in our MCMC approach. We take the prior distributions in chirp mass $\mathcal{M}_z$, reduced mass μ_z, polarization angle ψ, coalescence time t_c, and coalescence phase Φ_c to be *flat* over the region of sample space where the binary is detectable according to our selection procedure. More specifically, we choose

- $p^{(0)}(\mathcal{M}_z)$ = constant over the range $[1\,M_\odot,\ 2\,M_\odot]$ for NS-NS; and over the range $[2.5\,M_\odot,\ 4.9\,M_\odot]$ for NS-BH. (Note that the true chirp masses in the binaries' rest frames are $1.2\,M_\odot$ for NS-NS and $3.0\,M_\odot$ for NS-BH.)

- $p^{(0)}(\mu_z)$ = constant over the range $[0.3\,M_\odot,\ 2\,M_\odot]$ for NS-NS; and over the range $[0.5\,M_\odot,\ 3.5\,M_\odot]$ for NS-BH. (Note that the true reduced masses in the binaries' rest frames are $0.7\,M_\odot$ for NS-NS and $1.2\,M_\odot$ for NS-BH.)

- $p^{(0)}(\psi)$ = constant over the range $[0, \pi]$.

- $p^{(0)}(t_c)$ = constant over the range $[-100\,\mathrm{sec}, 100\,\mathrm{sec}]$. Since we assume that the coalescence time is close to the time of the SHB event, t_c is essentially the time offset between the system's final GWs and its SHB photons. We find that our choice of prior here is almost irrelevant, as long as the prior is flat and includes the true value. No matter how broad we choose the prior in t_c, our posterior PDF ends up very narrowly peaked around $\hat{t}_c$.

- $p^{(0)}(\Phi_c)$ = constant over the range $[0, 2\pi]$.

Given our assumed constant comoving density of SHBs, the prior distribution for luminosity distance scales as comoving volume over the range [0, 2 Gpc] for NS-NS binaries, and over the range [0, 5 Gpc] for NS-BH binaries. For our sample with isotropic inclination distribution, we put $p^{(0)}(\cos\iota)$ = constant over the range $[-1, 1]$. For our sample that assumes SHB collimation, our prior in $\cos\iota \equiv v$ is the same as the one that we used in our selection procedure discussed in the previous subsection:

$$\frac{dp^{(0)}}{dv}(v) \propto e^{-(1-v)^2/2\sigma_v^2}\ , \tag{53}$$

with $\sigma_v = 0.05$.

We then map out full distributions for each of our seven parameters, assessing the mean values [Eq. (48)] and the standard deviations [Eq. (49)]. We generate four chains which run in parallel on the CITA "Sunnyvale" Cluster. Each chain runs for a maximum of 10^7 steps; we find that the mean and median number of steps are $\sim 10^5$ and $\sim 10^4$, respectively. Each evaluation of the likelihood function takes ~ 0.3 seconds. We use the first 30% of a chain's sample states for "burn in," and thus discard that data. When generating our Markov chains, we use an initial Gaussian proposal distribution with a standard deviation similar to that expected for the posterior probability density:

$$Q(\boldsymbol{\theta}'; \boldsymbol{\theta}^{(n)}) = \frac{1}{(2\pi\sigma_T)^{N/2}} \exp[-\boldsymbol{\theta}^2/2\sigma_T^2], \tag{54}$$

where σ_T is the width of the trial Gaussian distribution and N is the number of dimensions of our parameter set (in our case, $N = 7$). However, in the case of parameters such as D_L and $\cos\iota$ that we expect to be strongly correlated, we use a much smaller proposal standard deviation. The proposal distribution is then updated using the covariance matrix of the last half of the samples.

Our chains start at random offset parameter values, drawn from Gaussians centered on the true parameter value. We assess convergence by testing whether the multiple chains have produced consistent parameter distributions. Following standard practice, we use the Gelman-Rubin convergence criterion, defining a sequence as "converged" if the statistic $R < 1.1$ on the last half of our samples; see Gelman & Rubin (1992) for more details. We use convergence as our stopping criterion.

Each simulation for every binary runs for an hour to forty-eight hours; the mean and median runtime are eight and three hours, respectively. CosmoMC features such as temperature annealing and sophisticated sampling methods have not been required in our simulations.

3.4. *The "averaged" posterior PDF*

Central to the procedure outlined above is the use of the datastream $\mathbf{s} = \mathbf{h}(\boldsymbol{\theta}) + \mathbf{n}$ which enters the likelihood function $\mathcal{L}_{\rm TOT}(\mathbf{s}|\boldsymbol{\theta})$. The resulting posterior PDF, and the parameters one infers, thus depend on the noise $\mathbf{n}$ which one uses, either via the noise in a particular measurement or (as in our case) by simulation (drawing randomly from the noise's PDF). In some cases, one will want to evaluate statistics that are in a well-defined sense "typical" given the average noise properties, rather than depending on a particular noise instance. Such averaging is appropriate, for example, when forecasting how well an instrument should be able to measure the properties of a source or process. As we discuss in more detail in the following Section, we have also found it is necessary to average when trying to compare our MCMC code's output with previously published work.

As explained in detail below, the averaged posterior PDF takes a remarkably simple form: It is just the "usual" posterior PDF, Eq. (39) with the noise $\mathbf{n}$ set to *zero*. We emphasize that this does not mean that one ignores noise when constructing the averaged PDF. For example, one still meaningfully relates the amplitude of the signal to the amplitude of typical rms noise in a detector by the average SNR, Eq. (45). As such, the averaged statistics will show an improvement in measurement accuracy as the SNR is increased.

To develop a useful notion of averaged posterior PDF, consider the hypothetical (and physically unrealistic) case in which we measure a signal using M different noise realizations for the same event. The joint posterior PDF for these measurements is

$$p_{\rm joint}(\boldsymbol{\theta}|\mathbf{s}_1, \mathbf{s}_2, \ldots \mathbf{s}_M) = \prod_{i=1}^{M} p(\boldsymbol{\theta}|\mathbf{s}_i)\ . \tag{55}$$

Let us define the "average" PDF as the geometric mean of the PDFs which describe these measurements:

$$p_{\rm ave}(\boldsymbol{\theta}|\mathbf{s}) \equiv p_{\rm joint}(\boldsymbol{\theta}|\mathbf{s}_1, \mathbf{s}_2, \ldots \mathbf{s}_M)^{1/M}\ . \tag{56}$$

Expanding this definition, we find

$$p_{\rm ave}(\boldsymbol{\theta}|\mathbf{s}) \equiv \prod_{i=1}^{M} [p(\boldsymbol{\theta}\,|\,\mathbf{s}_i)]^{1/M}$$

$$= \prod_{i=1}^{M} \left[\mathcal{N}_i \, p^{(0)}(\boldsymbol{\theta}) \mathcal{L}_{\rm TOT}(\mathbf{s}_i \,|\, \boldsymbol{\theta}) \right]^{1/M}$$
$$= \mathcal{N} \, p^{(0)}(\boldsymbol{\theta}) \prod_{i=1}^{M} \left[\mathcal{L}_{\rm TOT}(\mathbf{s}_i \,|\, \boldsymbol{\theta}) \right]^{1/M} , \quad (57)$$

where the subscript i denotes the ith noise realization in our set of M observations. The multi-observation likelihood can in turn be expanded as

$$\prod_{i=1}^{M} \left[\mathcal{L}_{\rm TOT}(\mathbf{s}_i|\boldsymbol{\theta}) \right]^{1/M} = \prod_{a} \prod_{i=1}^{M} \left[\mathcal{L}_a(s_{a,i} \,|\, \boldsymbol{\theta}) \right]^{1/M}$$
$$= \prod_{a} \prod_{i=1}^{M} e^{-\left(h_a(\boldsymbol{\theta}) - s_{a,i} \,\middle|\, h_a(\boldsymbol{\theta}) - s_{a,i} \right)/2M}$$
$$= \prod_{a} e^{-\left(h_a(\boldsymbol{\theta}) - h_a(\hat{\boldsymbol{\theta}}) \,\middle|\, h_a(\boldsymbol{\theta}) - h_a(\hat{\boldsymbol{\theta}}) \right)/2}$$
$$\times \prod_{i=1}^{M} \exp\left[\frac{1}{M} \left(n_{a,i} \middle| h_a(\boldsymbol{\theta}) - h_a(\hat{\boldsymbol{\theta}}) \right) \right]$$
$$\times \prod_{i=1}^{M} \exp\left[-\frac{1}{2M} \left(n_{a,i} \middle| n_{a,i} \right) \right] . \quad (58)$$

By taking M to be large, the products on the last two lines of Eq. (58) can be evaluated as follows:

$$\prod_{i=1}^{M} \exp\left[\frac{1}{M} \left(n_{a,i} \middle| h_a(\boldsymbol{\theta}) - h_a(\hat{\boldsymbol{\theta}}) \right) \right]$$
$$= \exp\left[\frac{1}{M} \sum_{i=1}^{M} \left(n_{a,i} \middle| h_a(\boldsymbol{\theta}) - h_a(\hat{\boldsymbol{\theta}}) \right) \right]$$
$$\simeq \exp\left[\left\langle \left(n_a \middle| h_a(\boldsymbol{\theta}) - h_a(\hat{\boldsymbol{\theta}}) \right) \right\rangle \right]$$
$$= 1 . \quad (59)$$

Here, $\langle \ldots \rangle$ denotes an ensemble average over noise realizations (cf. Sec. 3.1), and we have used the fact that our noise has zero mean. Similarly, we find

$$\prod_{i=1}^{M} \exp\left[-\frac{1}{2M} \left(n_{a,i} \middle| n_{a,i} \right) \right]$$
$$= \exp\left[-\frac{1}{2M} \sum_{i=1}^{M} \left(n_{a,i} \middle| n_{a,i} \right) \right]$$
$$\simeq \exp\left[-\frac{1}{2} \left\langle \left(n_a \middle| n_a \right) \right\rangle \right]$$
$$= e^{-1} . \quad (60)$$

The final equality uses $\langle (n_a|n_a) \rangle = 2$, which can be proved using using the noise properties (35), (36), and (37).

Putting all this together, we finally find

$$p_{\rm ave}(\boldsymbol{\theta}|\mathbf{s}) = \mathcal{N} p^0(\boldsymbol{\theta}) \prod_{a} e^{-\left(h_a(\boldsymbol{\theta}) - h_a(\hat{\boldsymbol{\theta}}) \,\middle|\, h_a(\boldsymbol{\theta}) - h_a(\hat{\boldsymbol{\theta}}) \right)/2} , \quad (61)$$

where we have absorbed the factor e^{-1} into the normalization $\mathcal{N}$. *The posterior PDF, averaged over noise realizations, is simply obtained by evaluating Eq. (39) with the noise* $\mathbf{n}$ *set to zero.*

4. RESULTS I: VALIDATION AND TESTING

We now validate and test our MCMC code against previously published results from CF94. In particular, we examine the posterior PDF for a particular NS-NS binary which was studied in detail in CF94. We also explore the dependence of distance measurement accuracies on the detector network and luminosity distance, focusing on the strong degeneracy that exists between $\cos\iota$ and D_L.

4.1. *Comparison with CF94*

Validation of our MCMC results requires comparing to work which goes beyond the Gaussian Fisher matrix approximation to the likelihood function. In Section IVD of CF94, Cutler & Flanagan investigate dominant effects that are non-linear in 1/SNR, and consequently, in the variable $1/D_L$. Their work shows that including such effects has a significant impact on the predicted distance measurement accuracies, particularly in the limit of low SNR. In particular, they find that Fisher-based estimates understate distance measurement errors for a network comprising the two LIGO detectors and Virgo.

Because they go beyond a Fisher matrix analysis, the results of CF94 are a useful comparison to our MCMC results. Their paper is also a useful touchstone in that they take a binary's source position to be known. It should be noted that their motivation for this assumption is different from ours: while we assume that the inspiral is associated with an SHB which determines the sky position, they argue that GWs on their own determine sky position precisely enough to break correlations between position and D_L. This argument is based on the "Marković approximation" (Marković 1993), which argues that timing information between different detectors fixes the sky position to sufficient accuracy.

Our approach is sufficently different from CF94 that we do not expect perfect agreement between our results. The most important difference is that we directly map out the posterior PDF and compute sample averages using Eqs. (48) and (49), for all of our parameters $\{\mathcal{M}_z, \mu_z, D_L, \cos\iota, \psi, t_c, \Phi_c\}$. In contrast, CF94 estimate measurement errors for only one parameter, D_L, using an approximate method based on an analytical Bayesian derivation of the marginalized PDF for D_L. Specifically, Cutler & Flanagan expand the exponential factor in Eq. (39) beyond second order in terms of some "best-fit" maxiumum likelihood parameters. Their approximation treats strong correlations between the parameters D_L and $\cos\iota$ that are non-linear in 1/SNR. However, other correlations between D_L and (ψ, ϕ_c) are only considered up to linear order. They obtain an analytical expression for the posterior PDF of the variables D_L and $\cos\iota$ in terms of their "best-fit" maximum-likelihood values $\tilde{D}_L$ and $\cos\tilde{\iota}$ [see Eq. (4.57) of CF94]. The marginalized 1-D posterior PDFs for D_L are then computed by numerically integrating over the parameter $\cos\iota$. We use the term 1-D marginalized PDF in parameter θ_i to refer to the distribution

$$p_{\rm marg}(\theta_i|\mathbf{s}) = \int \ldots \int p(\boldsymbol{\theta}|\mathbf{s}) d\theta_1 \ldots d\theta_{i-1} \, d\theta_{i+1} \ldots d\theta_N \quad (62)$$

where $p(\boldsymbol{\theta}|\mathbf{s})$ is the posterior PDF given by Eq. (39) and N is the number of dimensions of our parameter set (in

our case, $N = 7$).

In addition to this rather significant difference in our techniques, there are several relatively minor differences which also affect our comparison:

- We use the restricted 2PN waveform; CF94 use the leading "Newtonian, quadrupole" waveform (cf. the waveform we use for pedagogical purposes in Sec. 1.2). Since the distance is encoded in the waveform's amplitude, we do not expect that our use of a higher-order phase function will have a large impact. However, to avoid any easily circumvented mismatch, we adopt the Newtonian-quadrupole waveform for the purpose of our comparisons. Note that this waveform does not depend on the reduced mass μ; hence, *for the purpose of this comparison only*, our parameter space is reduced to 6 dimensions.

- We use the projected advanced sensitivity noise curve shown in Fig. (1); CF94 use an analytical form [their Eq. (2.1)[9]] based on the best-guess for what advanced sensitivity would achieve at the time of their analysis. Compared to the modern projected sensitivity, their curve underestimates the noise at middle frequencies ($\sim 40-150$ Hz) and overestimates it at high frequencies ($\gtrsim$ 200 Hz). We likewise adopt their noise curve for our comparison. Because of these differences, CF94 rather seriously overestimates the SNR for NS-NS inspiral. Using their noise curve, the average SNR for the binary presented in their Fig. 10 is 12.4[10]; using our up-to-date model for advanced LIGO, it is 5.8. As such, the reader should view the numbers in this section of our analysis as useful *only* for validation purposes.

- The two analyses use somewhat different prior distributions. As extensively discussed in Sec. 3.3, we set uniform priors on the chirp mass $\mathcal{M}_z$, on the time t_c and phase Φ_c at coalescence, and on the polarization phase ψ. For the purpose of this comparison, we will assume isotropic emission, so we also set a flat prior on the inclination $\cos\iota$. We assume our sources are uniformly distributed in constant comoving volume. However, we set a detection threshold which depends on the total network SNR, which effectively sets a joint prior on source inclination and distance. CF94 use a prior distribution only for the parameter set $\{D_L, \cos\iota, \psi, \Phi_c\}$ that is flat in polarization phase, coalescence phase, and inclination. [This smaller prior set follows from their use of the Marković approximation, which also argues that correlations between intrinsic and extrinsic parameters can be subsumed into correlations with the single intrinsic parameter Φ_c.] They assume a prior that is uniform in volume, but that cuts off the distribution at a distance $D_{\rm max} \simeq 6.5\,{\rm Gpc}$.

[9] Note that it is missing an overall factor of 1/5 (E. E. Flanagan, private communication).

[10] CF94 actually report the SNR to be 12.8. The slight discrepancy is due to rounding errors in the parameter r_0 in their Eq. (4.28). Adjusting to their preferred value (rather than computing r_0 directly) gives perfect agreement.

Our goal here is to reproduce the 1-D marginalized posterior PDF in D_L for the binary shown in Fig. 10 of CF94. For brevity, we refer to this system as the "CF binary." Each NS in the CF binary has $m_z = 1.4\,M_\odot$; it has a sky position $(\theta, \phi) = (50^\circ, 276^\circ)$; and the detector network comprises LIGO Hanford, LIGO Livingston and Virgo. CF94 report the "best-fit" maximum-likelihood values $(\tilde{D}_L, \cos\tilde{\iota}, \tilde{\Psi})$ to be $(432\,{\rm Mpc}, 0.31, 101.5^\circ)$, where $\Psi = \psi + \Delta\psi(\mathbf{n})$, and where $\Delta\psi(\mathbf{n})$ depends on the preferred basis of $\mathbf{e}^\times$ and $\mathbf{e}^\times$ set by the detector network [see Eqs. (4.23)–(4.25) of CF94[11]]. In their Bayesian framework, the GW signal has been generated and an experimenter has no access to the true parameters; they thus do not present the parameters' true values. In order to compare our distribution with theirs, we assume that $\hat{\boldsymbol{\theta}} = \tilde{\boldsymbol{\theta}}_{\rm ML}$ for the purpose of computing the likelihood function $\mathcal{L}(\boldsymbol{\theta}|\mathbf{s})$. This is a reasonable assumption in the limit in which priors are taken to be uniform over the relevant parameter space. As already mentioned, for this comparison we use their advanced detector noise curve and the Newtonian-quadrupole waveform. Finally, we interpret the solid curve in Fig. 10 of CF94 as the marginalized 1-D posterior PDF in D_L for an average of posterior PDFs of parameters (given an ensemble of many noisy observations for a particular event). We compute the average PDF as described in Sec. 3.4, and then marginalize over all parameters except D_L, as in Eq. (62).

The left hand panels of Fig. 6 show the resulting 1-D marginalized PDF in D_L and $\cos\iota$. Notice that its shape has a broad structure not dissimilar to the solid curve shown in Fig. 10 of CF94: The distribution has a small bump near $D_L \approx 460\,{\rm Mpc}$, a main peak at $D_L \approx 700\,{\rm Mpc}$, and extends out to roughly 1 Gigaparsec. Because of the broad shape, the Bayesian mean ($\tilde{D}_{L,{\rm BAYES}} = 694\,{\rm Mpc}$) is significantly different from both the true value ($\hat{D}_L = 432\,{\rm Mpc}$ in our calculation) and from the maximum likelihood ($\tilde{D}_{L,{\rm ML}} = 495\,{\rm Mpc}$). Note that, thanks to the marginalization, the peak of this curve does not coincide with the maximum likelihood. (Interestingly, we find a shape much closer to CF94 Fig. 10 if we use a prior in which binaries are uniformly distributed in distance, rather than uniformly distributed in volume.)

We further determine the 2-D marginalized posterior PDFs in D_L and $\cos\iota$ for the CF binary. Fig. 6 illustrates directly the very strong degeneracy between these parameters, as expected from the form of Eqs. (7) and (8), as well as from earlier works (e.g., Marković 1993, CF94). It's worth noting that, as CF94 comment, the binary they chose is measured particularly poorly. This is largely due to the fact that one polarization is measured by the network far better than the other, so that the D_L–$\cos\iota$ degeneracy remains relatively unbroken. This degeneracy is responsible for the characteristic tail to large D_L we find in the 1-D marginalized posterior PDF in D_L, $p(D_L|\mathbf{s})$, which we investigate further in the following section.

[11] Note that Eq. (4.25) of CF94 should read $\tan(4\Delta\psi) = 2\Theta_{+\times}/(\Theta_{++} - \Theta_{\times\times})$. In addition, $\tilde{\Psi} = 56.5^\circ$ should read $\tilde{\Psi} = 101.5^\circ$ under the caption of Fig. 10. (We have changed notation from $\bar{\psi}$ in CF94 to Ψ to avoid multiple accents on the best fit value.) We thank Éanna Flanagan for confirming these corrections.

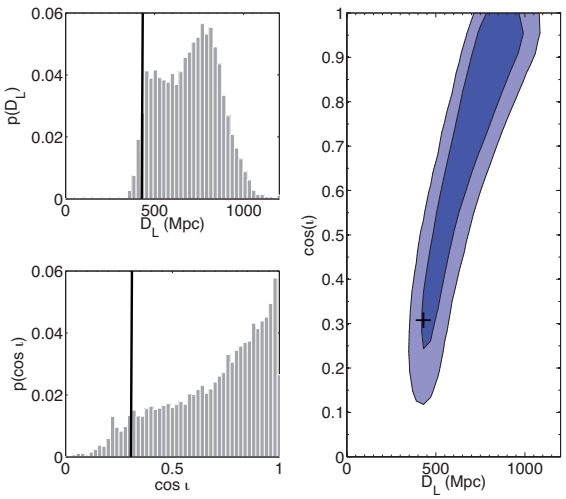

FIG. 6.— The 1-D and 2-D marginalized posterior PDFs for D_L and $\cos\iota$, averaged (as described in Sec. 3.4) over noise ensembles, for the "CF binary." The purpose of this calculation is to reproduce, as closely as possible, the non-Gaussian limit summarized in Fig. 10 of CF94. Top left-hand panel shows the 1-D marginalized posterior PDF in D_L (the true parameter value $\hat{D}_L = 432\,\mathrm{Mpc}$ is marked with a solid black line); bottom left-hand panel illustrates the 1-D marginalized posterior PDF in $\cos\iota$ (where the true value $\cos\hat{\iota} = 0.31$ is indicated with the solid black line). The right-hand panel shows the 2-D marginalized posterior PDF for the parameters D_L and $\cos\iota$, where the true values $\hat{D}_L$ and $\cos\hat{\iota}$ are marked with a cross. The contours around the dark and light blue areas indicate the 68 and 95% interval levels, respectively. Notice that the true values lie within the parameter interval defined by the contours of the 68% region. The Bayesian mean and rms measurement accuracies are (694.4 Mpc, 0.70) and (162 Mpc, 0.229) for (D_L, $\cos\iota$), respectively.

4.2. *Test 1: Varying luminosity distance and number of detectors*

We now examine how well we measure D_L as a function of both the distance to the CF binary and the properties of the GW detector network. Figures 7 and 8 show the 1-D and 2-D marginalized posterior PDFs in D_L and $\cos\iota$ for the CF binary at $\hat{D}_L = \{100, 200, 300, 400, 500, 600\}\,\mathrm{Mpc}$. For all these cases we keep the binary's sky position, inclination, and polarization angle fixed as in Sec. 4.1. The average SNRs we find for these six cases are (going from $\hat{D}_L = 100\,\mathrm{Mpc}$ to 600 Mpc) 53.6, 26.8, 17.9, 13.4, 10.7, and 8.9 (scaling, as expected, as $1/\hat{D}_L$). Interestingly, the marginalized PDFs for both distance and $\cos\iota$ shown in Figs. 7 and 8 have fairly Gaussian shapes for $\hat{D}_L = 100$ and 200 Mpc, but have very non-Gaussian shapes for $\hat{D}_L \geq 300\,\mathrm{Mpc}$. This can be considered "anecdotal" evidence that the Gaussian approximation for the posterior PDF breaks down at SNR $\lesssim 25$ or so, at least for this case. For lower SNR, the degeneracy between $\cos\iota$ and D_L becomes so severe that the 1-D errors on these parameters become quite large.

Next, we consider how measurement accuracy depends on properties of the GW detector network. Figure 9 shows the 1-D marginalized posterior PDFs in D_L for the CF binary with different networks. All "true" parameters are chosen exactly as described in Sec. 4.1. Adding detectors to the network does not substantially increase the total SNR; we increase the average total SNR from 12.4 to 14.6 (adding only AIGO), to 12.4 (adding only LCGT; its contribution is so small that the change is insignificant to the stated precision), or to 14.7 (adding both AIGO and LCGT). This change is not enough to counter the $D_L - \cos\iota$ degeneracy. While this degeneracy remains effective, the distance errors remain large and, in this case, biased. We remark also that the CF binary has a relatively small SNR for LCGT and Virgo: the average SNR in our detectors is 8.23 for LIGO-Hanford, 8.84 for LIGO-Livingston, 2.91 for Virgo, 8.71 for AIGO, and 1.1 for LCGT. This pathology of the CF binary is an example of a fairly general trend that we see. As we show in Sec. 5, even considering general binaries, randomly distributed on the sky and randomly oriented, we often find that SNR is quite low in one or more detectors.

4.3. *Test 2: Varying source inclination*

One of the prime results seen in our analysis of the CF binary is (as expected) a strong degeneracy between $\cos\iota$ and D_L. As Fig. 6 shows, this results in a tail to large values of D_L in the 1-D marginalized posterior PDF in D_L, $p(D_L|\mathbf{s})$, with a Bayes mean D_L for the CF binary of $\bar{D}_L = 694\,\mathrm{Mpc}$ (compared to the true CF value of 432 Mpc). Such a bias is naturally a great concern for using these sources as standard sirens, as well as GW measurements in general.

The CF binary has $\cos\hat{\iota} = 0.31$, meaning that it is nearly edge-on to the line of sight. Hypothesizing that the large tails may be due to its nearly edge-on nature, we consider a complementary binary that is nearly face on: We fix all of the parameters to those used for the CF binary, except for the inclination, which we take to be $\cos\hat{\iota} = 0.98$. We call this test case the "face-on" CF binary. By changing the inclination to a more nearly

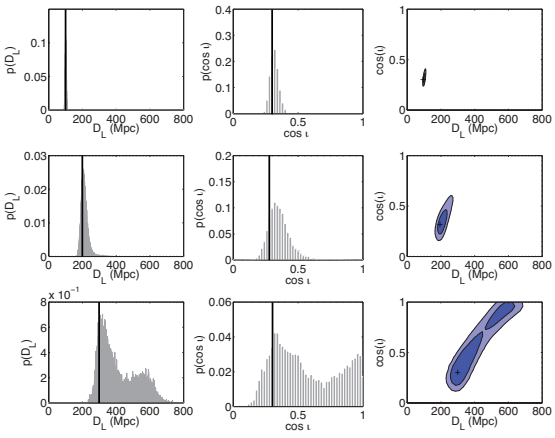

FIG. 7.— 1-D and 2-D marginalized PDFs for D_L and $\cos\iota$, averaged (as described in Sec. 3.4) over noise ensembles for the "CF binary" at different values of true luminosity distance $\hat{D}_L$: [100 Mpc, 200 Mpc, 300 Mpc] (top to bottom). The true values are marked with the solid black line, or a black cross in the 2-D case. The Bayesian means and rms measurement errors of the luminosity distance are [101.0 Mpc, 212.1 Mpc, 411.2 Mpc] and [3.6 Mpc, 21.4 Mpc, 110.0 Mpc], respectively. The corresponding Bayes mean and rms measurement errors for $\cos\iota$ are [0.317, 0.357, 0.562] and [0.033, 0.089, 0.247]. The dark and light contours in the 2-D marginalized PDF plots indicate the 68 and 95% interval levels, respectively. Notice that the true value always lies within the 68% contour region of the 2-D marginalized area at these luminosity distances.

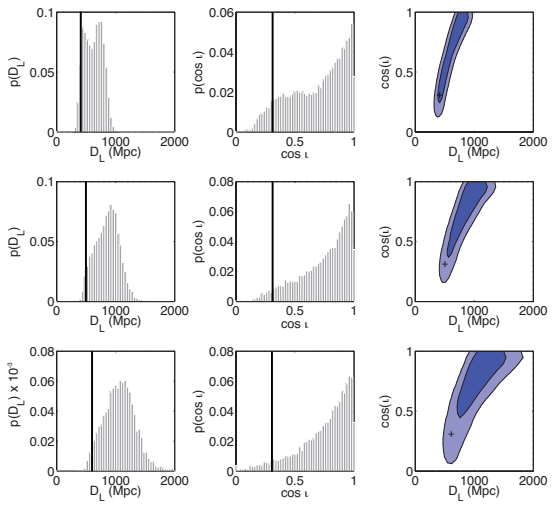

FIG. 8.— The 1-D and 2-D marginalized PDFs for D_L and $\cos\iota$, averaged (as described in Sec. 3.4) over noise ensembles for the "CF binary" at different values of the true luminosity distance $\hat{D}_L$: [400 Mpc, 500 Mpc, 600 Mpc] (top to bottom). The true values are marked with the solid black line, or a black cross in the 2-D case. The Bayesian means and rms measurement errors of the luminosity distance are [627.17 Mpc, 857.3 Mpc, 1068 Mpc] and [148.8 Mpc, 198.1 Mpc, 262.2 Mpc], respectively. The Bayes means and rms measurement errors for $\cos\iota$ are [0.686, 0.745, 0.746] and [0.237, 0.209, 0.218]. The dark and light contours in the 2-D marginalized PDF plots indicate the 68 and 95% interval levels, respectively. Notice that the true value lies within the 68% contour region of the 2-D marginalized area at a luminosity distance of 400 Mpc; for larger luminosity distances, the true value lies within the 95% contour region.

face-on situation, we substantially augment the measured SNR; the average SNR for the face-on CF binary measured by the LIGO/Virgo base network is 24.3 (as compared to 12.4 for the CF binary). We thus expect some improvement in the posterior PDF simply owing to the stronger signal.

Figure 10 shows the 1-D and 2-D marginalized posterior PDFs in D_L and $\cos\iota$. As we might have guessed, these distributions are complementary to those we found for the CF binary. In particular, we find that the peak of the 1-D marginalized posterior PDF in D_L is shifted to lower values in D_L, and the Bayes mean is much closer to the true value: $\tilde{D}_L = 376.3\,\mathrm{Mpc}$. Notice that the shape of the 1-D marginalized posterior PDF in $\cos\iota$ is abruptly cut off by the upper bound of the physical prior $\cos\iota \leq 1$. The Bayes mean for the inclination is $\cos\tilde{\iota} = 0.83$.

Just as we varied distance and detector network for the CF binary, we now do so for the face-on CF binary. Figures 11 and 12 show the 1-D marginalized posterior PDFs in D_L and $\cos\iota$ for the face-on CF binary for $\hat{D}_L = \{100, 200, 300, 400, 500, 600\}$ Mpc. The SNR in these cases is 105.0, 52.5, 34.0, 26.2, 20.1, and 17.5, respectively. Notice that the 1-D marginalized PDFs for $\cos\iota$ and D_L remain highly non-Gaussian, despite the high SNR. Once again, this is due to the hard cutoff in the PDF for $\cos\iota$. As long as a significant fraction of the PDF is excluded by the cutoff (i.e. as long as $1-\cos\iota \lesssim \sqrt{\Sigma^{\cos\iota\,\cos\iota}_{\mathrm{BAYES}}}$), then the 1-D marginalized 1-D PDFs for distance and inclination will be asymmetric and non-Gaussian. This implies that the Gaussian approximation and the Fisher matrix method will be unreliable. Similarly, Fig. 13 shows the 1-D marginalized posterior PDFs in D_L for the CF binary with different networks. Our results hence confirm this predicted degeneracy for face-on binaries irrespective of SNR.

4.4. *Discussion of validation and tests*

The main result from our validation tests is that the posterior PDFs we find have rather long tails, with strong correlations between $\cos\iota$ and D_L. Except for cases with very high SNR, the 1-D marginalized posterior PDF in $\cos\iota$ is rather broad. The Bayes mean for $\cos\iota$ thus typically suggests that a binary is at an intermediate inclination—underestimating $\cos\iota$ for nearly face-on binaries, and overestimating it for nearly edge-on binaries. Thanks to the strong $\cos\iota$–D_L degeneracy, we thus interpret (at fixed SNR) a nearly face-on binary as being closer than its true value, while interpreting a nearly edge-on binary as farther than its true value.

A bias in the measurement of D_L may be worrying for plans to use GW measurements as standard sirens. However, thus far we have only considered the measurement of single events. The joint constraints from measuring an ensemble of events can significantly reduce the bias, and improve the overall measurement. To see what happens when many events are observed, we consider the (physically unrealistic) case in which 80 NS-NS binaries are randomly placed on the sky with random orientation, but at fixed luminosity distance, $\hat{D}_L = 432\,\mathrm{Mpc}$, and measure these using our LIGO/Virgo baseline network. As we select binaries at a fixed distance on the sky, we now assume a uniform distance prior. We assemble the joint 1-D posterior likelihood for this particular

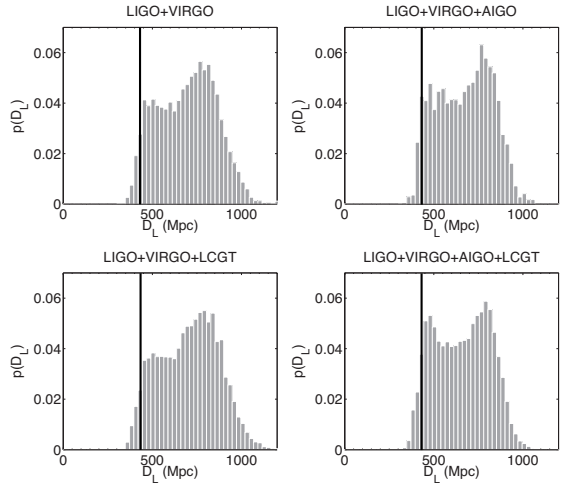

FIG. 9.— The 1-D marginalized PDF for D_L, averaged (as described in Sec. 3.4) over noise ensembles for the "CF binary" for different detector networks. The true values of $\hat{D}_L$ are marked with the solid black line. The lower right panel gives the case for the CF binary detected by LIGO, Virgo, AIGO, and LCGT. For this case, the Bayes mean and rms measurement error in D_L are [660.7 Mpc, 153.1 Mpc], respectively. The upper left is LIGO plus Virgo with [694.4 Mpc, 162.0 Mpc]; upper right is LIGO, Virgo, and AIGO, with [673.7 Mpc, 152.7 Mpc]; and lower left is LIGO, Virgo, and LCGT, with [706.1 Mpc, 164.6 Mpc]. Notice that the addition of detectors has little impact on improving measurements in D_L for this binary.

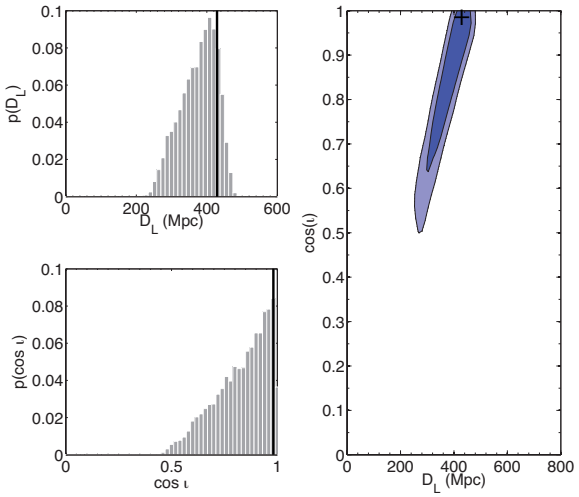

FIG. 10.— The 1-D and 2-D marginalized PDFs for D_L and $\cos\iota$, averaged (as described in Sec. 3.4) over noise ensembles for the "face-on" CF binary. The Bayesian mean and rms measurement accuracies are (376.3 Mpc, 0.83) and (51.3 Mpc, 0.12) for (D_L, $\cos\iota$), respectively. The top left-hand panel shows the 1-D marginalized posterior PDF in D_L (the true parameter value $\hat{D}_L = 432\,\mathrm{Mpc}$ is marked with a solid black line); the bottom left-hand panel illustrates the 1-D marginalized posterior PDF in $\cos\iota$ (where the true value $\cos\hat{\iota} = 0.98$ is indicated with the solid black line); and the right-hand panel demonstrates the 2-D marginalized posterior PDF for parameters D_L and $\cos\iota$, where the true values $\hat{D}_L$ and $\cos\hat{\iota}$ are marked with a cross. The dark and light contours indicate the 68 and 95% interval levels, respectively. Notice that the true values lie within the parameter interval defined by the contours of the 68% region.

distribution of events:

$$p_{\rm joint}(D_L) = \prod_{i=1}^{80} p_{\rm ave}(D_L|\mathbf{s}_i)\ , \tag{63}$$

where $p_{\rm ave}(D_L|\mathbf{s}_i)$ is the 1-D marginalized posterior PDF constructed from $p_{\rm ave}(\boldsymbol{\theta}|\mathbf{s}_i)$ for measurement i; the latter quantity having been averaged with respect to noise as discussed in Sec. 3.4. Figure 14 shows the joint 1-D marginalized posterior PDF for D_L. Notice that the bias is eliminated; in essence, averaging over many positions and orientations washes out the bias.

Though instructive, this limit is idealized. Real events will be distributed in distance, and a sample of eighty events is not expected in the near future. In the following section, we survey how well distances can be determined for realistic ensembles of GW-SHB events.

5. RESULTS II: SURVEY OF STANDARD SIRENS

We now examine how well various detector networks can measure an ensemble of canonical GW-SHB events. We consider NS-NS systems, with each neutron star having mass $m_z^{\rm NS} = (1+z)1.4\,M_\odot$; and NS-BH systems, with masses $m_z^{\rm NS} = (1+z)1.4\,M_\odot$ and $m_z^{\rm BH} = (1+z)10\,M_\odot$. We examine measurement by the four detector networks we have discussed (LIGO and Virgo; LIGO, Virgo, and AIGO; LIGO, Virgo, and LCGT; and LIGO, Virgo, AIGO, and LCGT). Finally, we consider both isotropic orientation and a beamed subsample, imagining that the gamma-rays from SHBs are beamed along the poles of the binary.

For this study we randomly choose events from our sample of *detected* NS-NS and NS-BH binaries (where the selection is detailed in Sec. 3.2). A rough estimate of the distance to which we can detect these events can be derived as follows. We set a total detector network threshold of 7.5, implying a threshold per detector of $7/\sqrt{5} = 3.4$ for a five detector network. Further averaging Eq. (45) over all sky positions and orientations yields (DHHJ06)

$$\left(\frac{S}{N}\right)_{a,\ \mathrm{sky-ave}} = \frac{8}{5}\sqrt{\frac{5}{96}}\frac{c}{D_L}\frac{1}{\pi^{2/3}}\left(\frac{G\mathcal{M}_z}{c^3}\right)^{5/6} \times \int_{f_{\rm low}}^{f_{\rm ISCO}} \frac{f^{-7/3}}{S_h(f)}df\ , \tag{64}$$

where the subscript "sky-ave" denotes that is the noise-averaged SNR for detector a, averaged over all sky positions and orientations. For a single detector threshold of 3.4, we find that a five detector network has an average range of about 600 Mpc for NS-NS events, and about 1200 Mpc for NS-BH events. If SHBs are associated with face-on binary inspiral, these numbers are increased by a factor $\sqrt{5/4} \simeq 1.12$ (DHHJ06).

If a constant comoving rate of 10 SHBs $\mathrm{Gpc}^{-3}\ \mathrm{yr}^{-1}$ is assumed (Nakar et al. 2006), we expect approximately 6 GW-SHB events per year for isotropically orientated NS-NS binary progenitors, and 44 SHBs per year for isotropically orientated NS-BH binary progenitors. If these events are face-on, the factor 1.12 increases the expected rate to 8 NS-NS and 57 NS-BH GW-SHB events per year. We stress that this is a rough approximation, since there are large uncertainties in the SHB event rate and redshift distribution.

In all cases we build our results by constructing the posterior distribution for an event given a unique noise realisation at each detector. We keep the noise realization, in a given detector and for a specific binary, constant as we add other detectors. This allows us to make

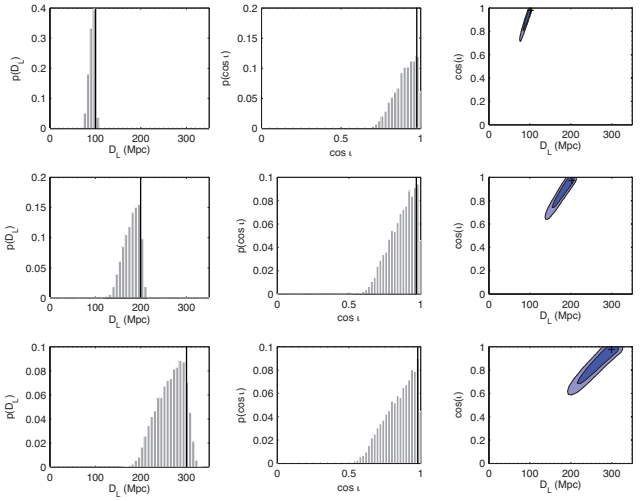

FIG. 11.— The 1-D marginalized PDFs for D_L and $\cos\iota$, averaged (as described in Sec. 3.4) over noise ensembles, for the face-on CF binary at different values of the true luminosity distance $\hat{D}_L$: [100 Mpc, 200 Mpc, 300 Mpc] (top to bottom). These values are marked with the solid black line. The Bayesian means and rms measurement errors of the luminosity distance are [92.4 Mpc, 179.4 Mpc, 264.5 Mpc] and [6.51 Mpc, 17.3 Mpc, 30.4 Mpc], respectively. The counterpart Bayes mean and rms measurement errors for $\cos\iota$ are [0.901, 0.869, 0.851] and [0.068, 0.092, 0.107]. The dark and light contours in the 2-D marginalized PDF plots indicate the 68 and 95% interval levels, respectively.

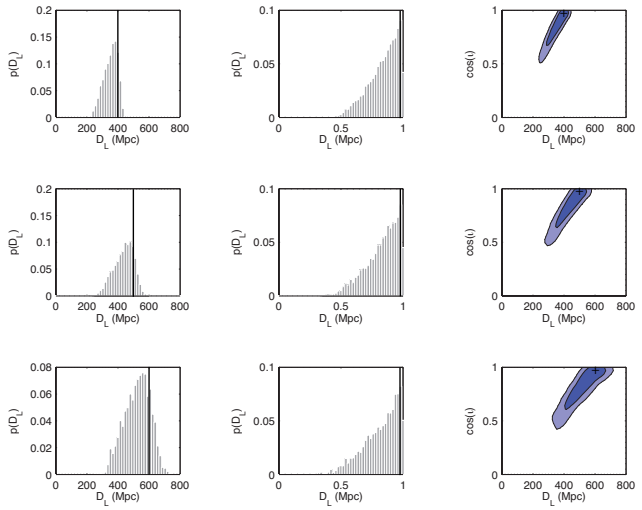

FIG. 12.— The 1-D and 2-D marginalized PDFs for D_L and $\cos\iota$, averaged (as described in Sec. 3.4) over noise ensembles, for the face-on CF binary at different values of the true luminosity distance $\hat{D}_L$: [400 Mpc, 500 Mpc, 600 Mpc] (top to bottom). The true values are marked with the solid black line. The Bayesian means and rms measurement errors for the luminosity distance are [352.2 Mpc, 437.8 Mpc, 524.5 Mpc] and [45.0 Mpc, 63.1 Mpc, 82.1 Mpc], respectively. The counterpart Bayes mean and rms measurement errors for $\cos\iota$ are [0.842, 0.830, 0.821] and [0.120, 0.131, 0.142]. The dark and light contours in the 2-D marginalized PDF plots indicate the 68 and 95% interval levels, respectively.

meaningful comparisons between the performance of difference detector networks.

5.1. *NS-NS binaries*

We begin with the case in which we have we have detected two hundred NS-NS binaries, either isotropically distributed in inclination angle or from our beamed subsample, using a network with all five detectors. Figure 15 shows scatter plots of the distance measurement accuracies for our unbeamed and beamed events, with each panel corresponding to a different detector network. The distance measurement error is defined as the ratio of the rms measurement error with the true value[12] $\hat{D}_L$:

$$\frac{\Delta D_L}{\hat{D}_L} = \frac{\sqrt{\Sigma^{D_L D_L}}}{\hat{D}_L} \,. \tag{65}$$

$\Sigma^{D_L D_L}$ is computed using (51). Although our small sample size precludes making definitive quantitative statements, we emphasize some general trends in by Fig. 15 which are particularly relevant to standard siren constraints on cosmological parameters:

- *The unbeamed total sample and the beamed subsample separate into two rather distinct distributions.* As anticipated, the beamed subsample improves measurement errors in D_L significantly, by greater than a factor of two or more. This is predominantly due to the associated beaming prior included in our analysis for these sources. The beaming prior constrains our inclination angle, $\cos\iota$, to $\sim 3\%$, thereby breaking the strong D_L–$\cos\iota$ degeneracy. By contrast, when no beaming prior is assumed, we find absolute errors of 0.1–0.3 in $\cos\iota$ for the majority of events; see Fig. 16. The strong D_L–$\cos\iota$ degeneracy then increases our distance errors. It's worth noting that a significant fraction of binaries randomly selected from our sample have $0.5 \lesssim |\cos\hat{\iota}| < 1$. As discussed in Sec. 3.2, this is due to the SNR selection criterion—at fixed distance, face-on binaries are louder and tend to be preferred.

- *When isotropic emission is assumed, we find a large scatter in distance measurement errors for all events, irrespective of network and true distance;* we find much less scatter when we assume a beaming prior. This is illustrated very clearly by the upper-right panel of Fig. 15. In that panel, we show the scatter of distance measurement error versus true distance for the LIGO, Virgo, AIGO detector network, comparing to the Fisher-matrix-derived linear scaling trend found in DHHJ06. For the unbeamed case (magenta points and line), our current results scatter around the linear trend; for the beamed case, most events lie fairly close to the trend. This demonstrates rather starkly the failure of Fisher methods to estimate distance measurement accuracy, especially when we cannot set a beaming prior.

- *Adding detectors to the network considerably increases the number of detected binaries, but does*

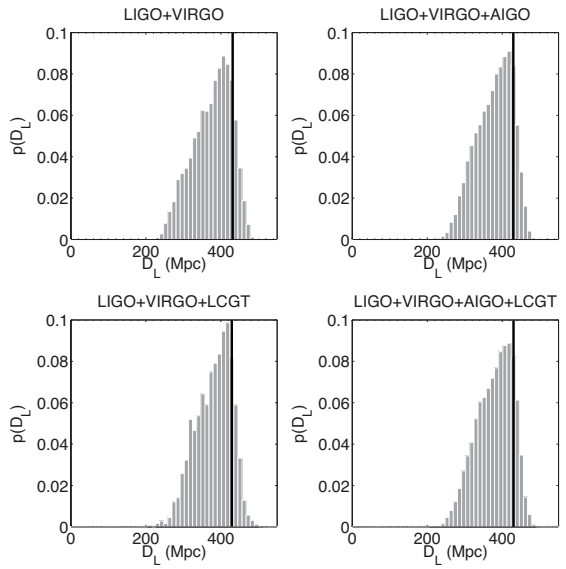

FIG. 13.— The 1-D marginalized PDF for D_L, averaged (as described in Sec. 3.4) over noise ensembles, for the face-on CF binary for different detector networks. The true values $\hat{D}_L$ are marked with the solid black line. The lower right-hand panel gives the case for the CF binary detected by LIGO, Virgo, AIGO, and LCGT. The Bayes mean and rms measurement error for D_L in this case is [379.2 Mpc, 48.0 Mpc], respectively. The upper left-hand panel is LIGO plus Virgo with [376.3 Mpc, 51.3 Mpc]; upper right is LIGO, Virgo, and AIGO, with [378.4 Mpc, 47.9 Mpc]; and lower left is LIGO, Virgo, and LCGT, with [380.6 Mpc, 47.6 Mpc]. As with the nearly edge-on CF binary, we again find that the additional detectors don't greatly improve measurements in D_L.

[12] Our definition differs from that given in CF94, their Eq. (4.62). There, the distance measurement is described as the ratio of the rms measurement error with the Bayes mean. We prefer to use Eq. (65) as we are interested primarily in the measurement error given a binary at its true luminosity distance.

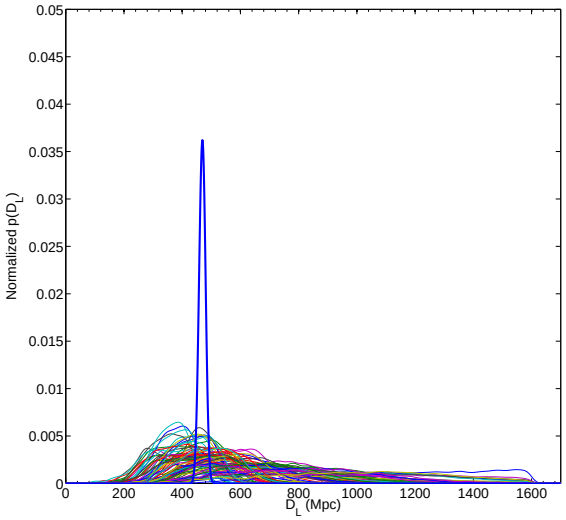

FIG. 14.— The 1-D marginalized PDFs for D_L for eighty NS-NS binaries randomly placed on the sky with random orientation, but at fixed luminosity distance, $\hat{D}_L = 432\,\mathrm{Mpc}$. We assume measurement using the LIGO and Virgo detectors. The thin colored lines show the normalized 1-D marginalized posterior PDFs for each NS-NS binary, where each event has been averaged (as described in Sec. 3.4) over all noise ensembles. The thick blue line illustrates the normalized joint 1-D PDF for D_L for the eighty events. The mean and standard deviation in the mean of the luminosity distance for eighty binaries are 468.1 Mpc and 10.8 Mpc, respectively. The standard deviation in the luminosity distance for each binary is approximately $10.8\,\mathrm{Mpc} \times \sqrt{80} \approx 96.6\,\mathrm{Mpc}$.

not significantly improve the accuracy with which those binaries are measured. The increase we see in the number of detected binaries is particularly significant for GW-SHB standard sirens. For instance, an important application is mapping out the posterior PDF for the Hubble constant H_0. As the number of events increases, the resulting joint posterior PDF in H_0 will become increasingly well constrained. Additional detectors also increases the distance to which binaries can be detected. This can be seen in Fig. 15: for the LIGO and Virgo network, our detected events extend to $\hat{D}_L \sim 600\,\mathrm{Mpc}$; the larger networks all go somewhat beyond this. Interestingly, networks which include the AIGO detector seem to reach somewhat farther out.

It is perhaps a bit disappointing that increasing the number of detectors in our network does not improve measurement accuracy. We believe this is due to two effects. First, a larger network tends to detect a larger number of weaker signals, and thus the additional binaries are poorly constrained. Second, the principle limitation to our measurement accuracy is the D_L–$\cos\iota$ degeneracy. A substantial improvement in distance measurement accuracy would require truly breaking this degeneracy (e.g., by applying the beaming prior).

5.2. *NS-BH binaries*

We now repeat the preceding analysis for NS-BH binaries. Figure 17 shows scatter plots of measurement accuracies for unbeamed and beamed NS-BH binaries. We find similar trends as in the NS-NS case:

- *The unbeamed and beamed samples separate into two distinct distributions.* Notice, however, that outliers exist in measurement errors at high D_L for several beamed events for all networks. This is not too surprising, given we expect beamed sources at higher luminosity distances and lower SNR. Such events are more likely to deviate from the linear relationship predicted by the Fisher matrix.

- *We see substantial scatter in distance measurement, particularly when isotropic emission is assumed.* As with NS-NS, the scatter is not as severe when we assume beaming, and in that case lies fairly close to a linear trend, as weould be predicted by a Fisher matrix. (Note that this trend is shallower in slope than for NS-NS binaries, thanks to the larger mass of the system.)

- *We do not see substantial improvement in distance measurement as we increase the detector network.* As with NS-NS, adding detectors does increase the range of the network; AIGO appears to particularly add events at large $\hat{D}_L$ (for both the isotropic and beamed samples). However, adding detectors does not break the fundamental D_L–$\cos\iota$ degeneracy. From our full posterior PDFs, we find absolute errors of 0.1–0.3 in $\cos\iota$, very much like in the NS-NS case.

6. SUMMARY DISCUSSION

In this analysis, we systematically study how well gravitational waves can be used to measure luminosity distances, under the assumption that short-hard gamma ray bursts are accompanied by binary inspiral. We examine two plausible compact binary SHB progenitors, and a variety of plausible detector network configurations. We emphasize that *we assume sky position is known.* We build on the previous study of DHHJ06, which used the so-called Gaussian approximation of the

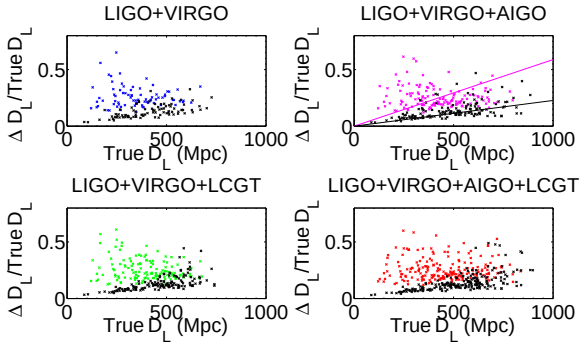

FIG. 15.— Distance measurement errors versus the true luminosity distance for our sample of NS-NS binaries. Colored points assume isotropic emission in our priors; black crosses assume our beaming prior. The lower right-hand panel shows the result for our "full" network (LIGO, Virgo, AIGO, and LCGT); in this case, 200 unbeamed and 200 beamed binaries are included. Upper left shows LIGO plus Virgo (84 unbeamed and 90 beamed); upper right is LIGO, Virgo, and AIGO (140 unbeamed and 138 beamed); lower left is LIGO, Virgo, and LCGT (141 unbeamed and 141 beamed). In the LIGO, Virgo, AIGO panel we also show the Fisher-matrix-derived linear scaling given in DHHJ06: $\Delta D_L/\hat{D}_L \simeq \hat{D}_L/(4.4\,\mathrm{Gpc})$ assuming beaming, and $\Delta D_L/\hat{D}_L \simeq \hat{D}_L/(1.7\,\mathrm{Gpc})$ for isotropic emission.

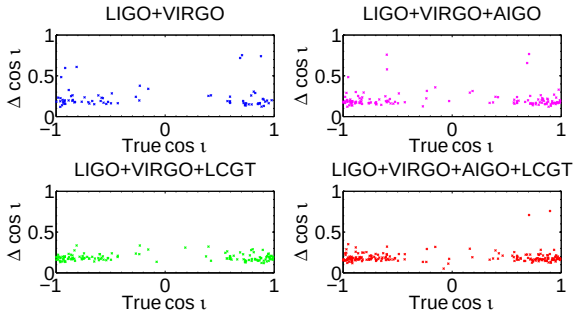

FIG. 16.— Inclination angle measurement errors versus true inclination angle for NS-NS binaries, assuming an isotropic orientation distribution. The lower right-hand panel is for the "full" LIGO, Virgo, AIGO, LCGT network and includes 200 binaries. Upper left is LIGO plus Virgo (81 binaries); upper right is LIGO, Virgo, and AIGO (140 binaries); and lower left is LIGO, Virgo, and LCGT (141 binaries).

posterior PDF. This approximation is known to work well for large SNR; the limits of its validity, however, are not well understood. Since the SNR of events measured by ground-based detectors is likely to be of order 10, we are concerned that the Gaussian limit may not be valid. We examine the posterior PDF for the parameters of measured events using Markov Chain Monte Carlo techniques, which do not rely on this approximation. We also introduce a well-defined posterior PDF that is averaged over all noise realizations and does not depend solely on a particular noise instance. Such a quantity is especially useful when we wish to predict how well a detector should be able to measure the properties of a source or process.

We find that the Gaussian description of the likelihood substantially underestimates distance measurement errors. We also find that the main limitation for individual standard siren measurements is the strong degeneracy between distance to the binary and its inclination to the line of sight. Adding detectors to a network does not substantially improve the distance measurement for a given single event. When we assume that the SHB is isotropic (so that we cannot infer anything about the source's inclination given the GRB measurement), we find that Fisher matrix estimates of distance errors are very inaccurate. Our distributions show very large scatter about Fisher-based predictions.

The situation is improved dramatically if we can assume that SHBs are collimated, thereby giving us a prior on the orientation of the progenitor binary. By assuming that the GRBs are preferentially emitted into an opening angle of roughly 25°, we find that the distance–inclination correlation is strongly broken. In this case, the Fisher matrix estimates are much more reasonable, giving a fairly good sense of the trend with which distances are determined (albeit with a moderate amount of scatter about that trend). This illustrates the importance of incorporating prior knowledge of our parameters into our measurement.

Our distance measurement results are summarized by Fig. 15 (for NS-NS SHB progenitors) and Fig. 17 (for NS-BH). Assuming isotropy, we find the distance to NS-NS binaries is measured with a fractional error of roughly 20–60%, with most events in our distribution clustered near 20–30%. Beaming improves this by roughly a factor of two, and elminates most of the high error events from our sample. Similar results are found for NS-BH events, with perhaps an improvement of 10% or so.

At first blush these findings appear disheartening with regards to standard siren measurements. We emphasize, however, that these results describe the outcome of *individual siren measurements*. When these measurements are used as cosmological probes, we will be interested in constructing the joint distribution, following observation of N GW-SHB events. As Fig. 14 illustrates, the joint distribution found by combining many individual measurements becomes increasingly well constrained, washing out the scatter we see on an event-by-event basis. Indeed, preliminary studies show that our ability to constrain H_0 improves quite sharply as the number of measured binaries is increased. In our most pessimistic scenario (the SHB is assumed to be a NS-NS binary, with no prior on inclination, and measured by the baseline LIGO-Virgo network), we find that H_0 can be measured with $\sim 13\%$ fractional error with $N = 4$, improving to $\sim 5\%$ for $N = 15$. We stress that this measurement is based on absolute determinations of distance using standard sirens. (Details of this analysis will be presented in a companion paper, currently in preparation.)

Increasing the number of measured events will thus be crucial for getting cosmologically interesting measurements. To this end, it is important to note that increasing the number of detectors in our network enables a considerable increase in the number of detected binaries. This is due to both improvement in sky coverage, and in the total detection volume. Going from a network which includes all four detectors (LIGO, Virgo, AIGO, and LCGT) to our baseline network of just LIGO and Virgo entails a $\sim 50\%$ reduction in the number of detected binaries. Eliminating just one of the proposed detectors (AIGO or LCGT) leaves us with $\sim 75\%$ of the original detected sample. Interestingly, we find that AIGO shows a marked improvement in certain areas of the sky, resulting in a smaller scatter of measurement accuracies for an ensemble of events at low SNR.

Assuming an event rate density of 10 SHBs per year per Gpc^3 (Nakar et al. 2006), we expect to measure 6 NS-NS events per year or 44 NS-BH events per year using a network with all five detectors. For the LIGO-Virgo network, we expect half this rate; for the LIGO-Virgo-

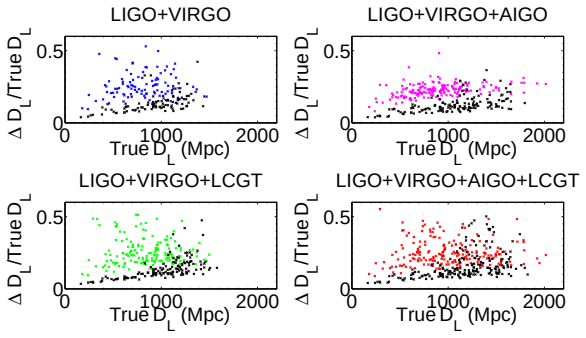

FIG. 17.— Distance measurement errors versus the true luminosity distance for our sample of NS-BH binaries. Colored points assume isotropic emission; black crosses use our beaming prior. The lower right-hand panel shows the sample detected by our "full" network (LIGO, Virgo, AIGO, LCGT) and includes 250 unbeamed and 200 beamed binaries. Upper left is LIGO plus Virgo (117 unbeamed and 98 beamed); upper right is LIGO, Virgo, and AIGO (180 unbeamed and 147 beamed); and lower left is LIGO, Virgo, and LCGT (179 unbeamed and 144 beamed).

AIGO or LIGO-Virgo-LCGT network, we expect 3/4 of this rate. If SHB collimation can be assumed, the rate is further augmented by a factor of 1.12. At this rate, we find that *one* year of observation should be enough to measure H_0 to an accuracy of $\sim 1\%$ if SHBs are dominated by beamed NS-BH binaries using the "full" network of LIGO, Virgo, AIGO, and LCGT—admittedly, our most optimistic scenario. A general trend we see is a network of five detectors (as opposed to our baseline LIGO-Virgo network of three detectors) increases measurement accuracy in H_0 by a factor of one and a half; assuming that the SHB progenitor is a NS-BH binary improves measurement accuracies by a factor of four or greater. Errors in H_0 are seen to improve by a factor of at least two when we assume SHB collimation.

Aside from exploring the cosmological consequences of these results, several other issues merit careful future analysis. One general result we found is the importance that prior distributions have on our final posterior PDF. We plan to examine this in some detail, checking which parameters particularly influence our final result, and ascertaining what uncertainties can be ascribed to our inability to set priors on these parameters. It may be possible to mitigate the influence of the D_L–$\cos\iota$ degeneracy by setting a distance prior that requires our inferred distance to be consistent with the SHB's observed redshift.

Another important issue is that of systematic errors in binary modeling. We have used the second-post-Newtonian description of a binary's GWs in our analysis; and, we have ignored all but the leading quadrupole harmonic of the waves (the so-called "restricted" post-Newtonian waveform). Our suspicion is that a more complete post-Newtonian description of the phase would have little impact on our results, since such effects are not likely to have an impact on the all-important D_L–$\cos\iota$ degeneracy. In principle, including additional (non-quadrupole) harmonics could have an impact on this degeneracy, since these other harmonics encode different information about the inclination angle ι. In practice, we expect that they won't have much effect on GW-SHB measurements, since these harmonics are measured with very low SNR (the strongest harmonic is roughly a factor of 10 smaller in amplitude than the quadrupole). It shouldn't be too difficult to test this, however; given how important this degeneracy has proven to be, it could be a worthwhile exercise.

As discussed previously, we confine our analysis to the inspiral part of the waveform. Inspiral waves are terminated at the presumed innermost stable circular orbit frequency, $f_{\rm ISCO} = (6^{3/2}\pi M_z)$. For NS-NS binaries, $f_{\rm ISCO} \simeq 1600\,{\rm Hz}$. At this frequency, detectors have fairly poor sensitivity, and we are thus confident that terminating the waves has little impact on our results for NS-NS systems. However, for our assumed NS-BH binaries, $f_{\rm ISCO} \simeq 400\,{\rm Hz}$. Detectors have rather good sensitivity in this band, so it may be quite important to improve our model for the waves' termination in this case.

Perhaps the most important follow-up would be to include the impact of spin. Although the impact of neutron star spin is likely to be small, it may not be negligible; and, for NS-BH systems, the impact of the black hole's spin could be significant. Spin induces precessions in the binary which can make the orientation of the orbit, $\hat{\mathbf{L}}$, dynamical. That in turn makes the observed inclination dynamical, which can break the D_L–$\cos\iota$ degeneracy. Van der Sluys et al. (2008) have already shown that spin precession physics vastly improves the ability of ground-based detectors to determine a source's position on the sky; we are confident that a similar analysis which assumes sky position will find that measurements of source distance and inclination can likewise be improved.

It is a pleasure to acknowledge useful discussions with Yoicho Aso, Duncan Brown, Curt Cutler, Jean-Michel Désert, Éanna Flanagan, Zhiqi Huang, Ryan Lang, Antony Lewis, Nergis Mavalvala, Szabolcs Márka, Phil Marshall, Cole Miller, Peng Oh, Ed Porter, Alexander Shirokov, and Pascal Vaudrevange. We are grateful to Neil Cornish in particular for early guidance on the development of our MCMC code. We also are grateful for the hospitality of the Kavli Institute for Theoretical Physics at UC Santa Barbara, and to the Aspen Center for Physics, where portions of the work described here were formulated. Computations were performed using the Sunnyvale computing cluster at the Canadian Institute for Theoretical Astrophysics, which is funded by the Canadian Foundation for Innovation. SAH is supported by NSF Grant PHY-0449884, and also gratefully acknowledges the support of the MIT Class of 1956 Career Development Fund.

REFERENCES

Althouse, W., Jones, L., & Lazzerini, A. 1998, Technical Report No. LIGO-T980044-08

Anderson, W. G., Brady, P. R., Creighton, J. D., & Flanagan, É. É. 2001, Phys. Rev. D, 63, 042003

Berger, E., Fox, D. B., Price, P. A., Nakar, E., Gal-Yam, A., Holz, D. E., Schmidt, B. P., Cucchiara, A., Cenko, S. B., Kulkarni, S. R., Soderberg, A. M., Frail, D. A., Penprase, B. E., Rau, A., Ofek, E., Burnell, S. J. B., Cameron, P. B., Cowie, L. L., Dopita, M. A., Hook, I., Peterson, B. A., Podsiadlowski, P., Roth, K. C., Rutledge, R. E., Sheppard, S. S., & Songaila, A. 2007, ApJ, 664, 1000

Blair, D. G., Barriga, P., Brooks, A. F., Charlton, P., Coward, D., Dumas, J.-C., Fan, Y., Galloway, D., Gras, S., Hosken, D. J., Howell, E., Hughes, S., Ju, L., McClelland, D. E., Melatos, A., Miao, H., Munch, J., Scott, S. M., Slagmolen, B. J. J., Veitch, P. J., Wen, L., Webb, J. K., Wolley, A., Yan, Z., & Zhao, C. 2008, Journal of Physics Conference Series, 122, 012001

Blanchet, L. 2006, Living Reviews in Relativity, 9, 4

Blanchet, L., Damour, T., Esposito-Farèse, G., & Iyer, B. R. 2004, Physical Review Letters, 93, 091101

Blanchet, L., Damour, T., Iyer, B. R., Will, C. M., & Wiseman, A. G. 1995, Phys. Rev. Lett., 74, 3515

Blanchet, L., Faye, G., Iyer, B. R., & Joguet, B. 2002a, Phys. Rev. D, 65, 061501

Blanchet, L., Iyer, B. R., & Joguet, B. 2002b, Phys. Rev. D, 65, 064005

Bloom, J. S., Holz, D. E., Hughes, S. A., Menou, K., Adams, A., Anderson, S. F., Becker, A., Bower, G. C., Brandt, N., Cobb, B., Cook, K., Corsi, A., Covino, S., Fox, D., Fruchter, A., Fryer, C., Grindlay, J., Hartmann, D., Haiman, Z., Kocsis, B., Jones, L., Loeb, A., Marka, S., Metzger, B., Nakar, E., Nissanke, S., Perley, D. A., Piran, T., Poznanski, D., Prince, T., Schnittman, J., Soderberg, A., Strauss, M., Shawhan, P. S., Shoemaker, D. H., Sievers, J., Stubbs, C., Tagliaferri, G., Ubertini, P., & Wozniak, P. 2009, ArXiv e-prints

Buonanno, A., & Damour, T. 1999, Phys. Rev. D, 59, 084006

Buonanno, A., Pan, Y., Pfeiffer, H. P., Scheel, M. A., Buchman, L. T., & Kidder, L. E. 2009, ArXiv e-prints
Burrows, D. N., Grupe, D., Capalbi, M., Panaitescu, A., Patel, S. K., Kouveliotou, C., Zhang, B., Mészáros, P., Chincarini, G., Gehrels, N., & Wijers, R. A. M. 2006, ApJ, 653, 468
Cavalier, F., Barsuglia, M., Bizouard, M.-A., Brisson, V., Clapson, A.-C., Davier, M., Hello, P., Kreckelbergh, S., Leroy, N., & Varvella, M. 2006, Phys. Rev. D, 74, 082004
Christensen, N., Meyer, R., & Libson, A. 2004, Class. Quantum Grav., 21, 317
Cook, G. B., Shapiro, S. L., & Teukolsky, S. A. 1994, ApJ, 424, 823
Cornish, N. J., & Porter, E. K. 2007, Classical and Quantum Gravity, 24, 5729
Cutler, C. 1998, Phys. Rev. D, 57, 7089
Cutler, C., & Flanagan, É. E. 1994, Phys. Rev. D, 49, 2658, referred to in the text as CF94
Cutler, C., & Vallisneri, M. 2007, Phys. Rev. D, 76, 104018
Dalal, N., Holz, D. E., Hughes, S. A., & Jain, B. 2006, Phys. Rev. D, 74, 063006, referred to in the text as DHHJ06
Droz, S., Knapp, D. J., Poisson, E., & Owen, B. J. 1999, Phys. Rev. D, 59, 124016
Dunkley, J., Komatsu, E., Nolta, M. R., Spergel, D. N., Larson, D., Hinshaw, G., Page, L., Bennett, C. L., Gold, B., Jarosik, N., Weiland, J. L., Halpern, M., Hill, R. S., Kogut, A., Limon, M., Meyer, S. S., Tucker, G. S., Wollack, E., & Wright, E. L. 2009, ApJS, 180, 306
Eichler, D., Livio, M., Piran, T., & Schramm, D. N. 1989, Nature, 340, 126
Etienne, Z. B., Faber, J. A., Liu, Y. T., Shapiro, S. L., Taniguchi, K., & Baumgarte, T. W. 2008, Phys. Rev. D, 77, 084002
Finn, L. S. 1992, Phys. Rev. D, 46, 5236
Finn, L. S., & Chernoff, D. F. 1993, Phys. Rev. D, 47, 2198
Ford, E. B. 2005, AJ, 129, 1706
Fox, D. B., Frail, D. A., Price, P. A., Kulkarni, S. R., Berger, E., Piran, T., Soderberg, A. M., Cenko, S. B., Cameron, P. B., Gal-Yam, A., Kasliwal, M. M., Moon, D.-S., Harrison, F. A., Nakar, E., Schmidt, B. P., Penprase, B., Chevalier, R. A., Kumar, P., Roth, K., Watson, D., Lee, B. L., Shectman, S., Phillips, M. M., Roth, M., McCarthy, P. J., Rauch, M., Cowie, L., Peterson, B. A., Rich, J., Kawai, N., Aoki, K., Kosugi, G., Totani, T., Park, H.-S., MacFadyen, A., & Hurley, K. C. 2005, Nature, 437, 845
Gelman, A., & Rubin, D. B. 1992, Statistical Science, 4, 457
Gilks, W. R., Richardson, S., & Spiegelhalter, D. J. 1996, Markov Chain Monte Carlo in Practice (Markov Chain Monte Carlo in Practice, by W. R. Gilks, S. Richardson and D. J. Spiegelhalter, pp. 486. ISBN 0412055511. London, UK: Chapman and Hall, 1996.)
Grupe, D., Burrows, D. N., Patel, S. K., Kouveliotou, C., Zhang, B., Mészáros, P., Wijers, R. A. M., & Gehrels, N. 2006, ApJ, 653, 462
Hastings, W. K. 1970, Biometrika, 57, 97
Hogg, D. W. 1999, ArXiv Astrophysics e-prints
Kochanek, C. S., & Piran, T. 1993, ApJ, 417, L17
Komatsu, E., Dunkley, J., Nolta, M. R., Bennett, C. L., Gold, B., Hinshaw, G., Jarosik, N., Larson, D., Limon, M., Page, L., Spergel, D. N., Halpern, M., Hill, R. S., Kogut, A., Meyer, S. S., Tucker, G. S., Weiland, J. L., Wollack, E., & Wright, E. L. 2009, ApJS, 180, 330
Kopparapu, R. K., Hanna, C., Kalogera, V., O'Shaughnessy, R., González, G., Brady, P. R., & Fairhurst, S. 2008, ApJ, 675, 1459
Królak, A., Lobo, J. A., & Meers, B. J. 1993, Phys. Rev. D, 48, 3451
Kulkarni, S., & Kasliwal, M. M. 2009, ArXiv e-prints
Lang, R. N., & Hughes, S. A. 2006, Phys. Rev. D, 74, 122001
Lewis, A., & Bridle, S. 2002, Phys. Rev. D, 66, 103511
Loredo, T. J. 1989, in Bulletin of the American Astronomical Society, Vol. 21, Bulletin of the American Astronomical Society, 1136
MacKay, D. J. C. 2003, Information Theory, Inference and Learning Algorithms (Information Theory, Inference and Learning Algorithms, by David J. C. MacKay, pp. 640. ISBN 0521642981. Cambridge, UK: Cambridge University Press, October 2003.)
Marković, D. 1993, Phys. Rev. D, 48, 4738
Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. 1953, J. Chem. Phys., 21, 1087
Nakar, E., Gal-Yam, A., & Fox, D. B. 2006, Astrophys. J., 650, 281
Neal, R. M. 1993, Technical Report CRG-TR-93-1
Owen, B. J. 1996, Phys. Rev. D, 53, 6749
Peebles, P. J. E. 1993, Principles of physical cosmology (Princeton Series in Physics, Princeton, NJ: Princeton University Press, —c1993)
Perley, D. A., Metzger, B. D., Granot, J., Butler, N. R., Sakamoto, T., Ramirez-Ruiz, E., Levan, A. J., Bloom, J. S., Miller, A. A., Bunker, A., Chen, H.-W., Filippenko, A. V., Gehrels, N., Glazebrook, K., Hall, P., Hurley, K. C., Kocevski, D., Li, W., Lopez, S., Norris, J., Piro, A. L., Poznanski, D., Prochaska, J. X., Quataert, E., & Tanvir, N. 2008, ArXiv e-prints
Phinney, E. S. 2009, ArXiv e-prints
Poisson, E., & Will, C. M. 1995, Phys. Rev. D, 52, 848
Porter, E. K., & Cornish, N. J. 2008, Phys. Rev. D, 78, 064005
Pretorius, F. 2005, Phys. Rev. Lett., 95, 121101
Röver, C., Meyer, R., & Christensen, N. 2007, Phys. Rev. D, 75, 062004
Schutz, B. F. 1986, Nature, 323, 310
Shibata, M., & Uryū, K. 2006, Phys. Rev. D, 74, 121503
Sievers, J. L., Mason, B. S., Weintraub, L., Achermann, C., Altamirano, P., Bond, J. R., Bronfman, L., Bustos, R., Contaldi, C., Dickinson, C., Jones, M. E., May, J., Myers, S. T., Oyarce, N., Padin, S., Pearson, T. J., Pospieszalski, M., Readhead, A. C. S., Reeves, R., Shepherd, M. C., Taylor, A. C., & Torres, S. 2009, ArXiv e-prints
Soderberg, A. M., Berger, E., Kasliwal, M., Frail, D. A., Price, P. A., Schmidt, B. P., Kulkarni, S. R., Fox, D. B., Cenko, S. B., Gal-Yam, A., Nakar, E., & Roth, K. C. 2006, ApJ, 650, 261
Stroeer, A., Gair, J., & Vecchio, A. 2006, in American Institute of Physics Conference Series, Vol. 873, Laser Interferometer Space Antenna: 6th International LISA Symposium, ed. S. M. Merkovitz & J. C. Livas, 444–451
Sylvestre, J. 2004, Class. Quantum Grav., 21, 775
Thorne, K. S. 1987, in 300 Years of Gravitation, ed. S. W. Hawking & W. Israel (Cambridge: Cambridge University Press), 330–458
van der Sluys, M. V., Röver, C., Stroeer, A., Raymond, V., Mandel, I., Christensen, N., Kalogera, V., Meyer, R., & Vecchio, A. 2008, ApJ, 688, L61
Wickham, E. D. L., Stroeer, A., & Vecchio, A. 2006, Classical and Quantum Gravity, 23, 819
Winn, J. N., Holman, M. J., Bakos, G. Á., Pál, A., Johnson, J. A., Williams, P. K. G., Shporer, A., Mazeh, T., Fernandez, J., Latham, D. W., & Gillon, M. 2007, AJ, 134, 1707