

Search for gravitational waves from galactic and extra-galactic binary neutron stars

- B. Abbott,¹² R. Abbott,¹⁵ R. Adhikari,¹³ A. Ageev,^{20,27} B. Allen,³⁹ R. Amin,³⁴ S. B. Anderson,¹² W. G. Anderson,²⁹ M. Araya,¹² H. Armandula,¹² M. Ashley,²⁸ F. Asiri,^{12,a} P. Aufmuth,³¹ C. Aulbert,¹ S. Babak,⁷ R. Balasubramanian,⁷ S. Ballmer,¹³ B. C. Barish,¹² C. Barker,¹⁴ D. Barker,¹⁴ M. Barnes,^{12,b} B. Barr,³⁵ M. A. Barton,¹² K. Bayer,¹³ R. Beausoleil,^{26,c} K. Belczynski,²³ R. Bennett,^{35,d} S. J. Berukoff,^{1,e} J. Betzwieser,¹³ B. Bhawal,¹² I. A. Bilenko,²⁰ G. Billingsley,¹² E. Black,¹² K. Blackburn,¹² L. Blackburn,¹³ B. Bland,¹⁴ B. Bochner,^{13,f} L. Bogue,¹² R. Bork,¹² S. Bose,⁴⁰ P. R. Brady,³⁹ V. B. Braginsky,²⁰ J. E. Brau,³⁷ D. A. Brown,¹² A. Bullington,²⁶ A. Bunkowski,^{2,31} A. Buonanno,^{6,g} R. Burgess,¹³ D. Busby,¹² W. E. Butler,³⁸ R. L. Byer,²⁶ L. Cadonati,¹³ G. Cagnoli,³⁵ J. B. Camp,²¹ C. A. Cantley,³⁵ L. Cardenas,¹² K. Carter,¹⁵ M. M. Casey,³⁵ J. Castiglione,³⁴ A. Chandler,¹² J. Chapsky,^{12,b} P. Charlton,^{12,h} S. Chatterji,¹³ S. Chelkowski,^{2,31} Y. Chen,⁶ V. Chickarmane,^{16,i} D. Chin,³⁶ N. Christensen,⁸ D. Churches,⁷ T. Cokelaer,⁷ C. Colacino,³³ R. Coldwell,³⁴ M. Coles,^{15,j} D. Cook,¹⁴ T. Corbitt,¹³ D. Coyne,¹² J. D. E. Creighton,³⁹ T. D. Creighton,¹² D. R. M. Crooks,³⁵ P. Csatorday,¹³ B. J. Cusack,³ C. Cutler,¹ E. D'Ambrosio,¹² K. Danzmann,^{2,31} E. Daw,^{16,k} D. DeBra,²⁶ T. Delker,^{1,34} V. Dergachev,³⁶ R. DeSalvo,¹² S. Dhurandhar,¹¹ A. Di Credico,²⁷ M. Díaz,²⁹ H. Ding,¹² R. W. P. Drever,⁴ R. J. Dupuis,³⁵ J. A. Edlund,^{12,b} P. Ehrens,¹² E. J. Elliffe,³⁵ T. Etzel,¹² M. Evans,¹² T. Evans,¹⁵ S. Fairhurst,³⁹ C. Fallnich,³¹ D. Farnham,¹² M. M. Fejer,²⁶ T. Findley,²⁵ M. Fine,¹² L. S. Finn,²⁸ K. Y. Franzen,³⁴ A. Freise,^{2,m} R. Frey,³⁷ P. Fritschel,¹³ V. V. Frolov,¹⁵ M. Fyffe,¹⁵ K. S. Ganezer,⁵ J. Garofoli,¹⁴ J. A. Giaime,¹⁶ A. Gillespie,^{12,n} K. Goda,¹³ G. González,¹⁶ S. Goßler,³¹ P. Grandclément,^{23,o} A. Grant,³⁵ C. Gray,¹⁴ A. M. Gretarsson,¹⁵ D. Grimmitt,¹² H. Grote,² S. Grunewald,¹ M. Guenther,¹⁴ E. Gustafson,^{26,p} R. Gustafson,³⁶ W. O. Hamilton,¹⁶ M. Hammond,¹⁵ J. Hanson,¹⁵ C. Hardham,²⁶ J. Harms,¹⁹ G. Harry,¹³ A. Hartunian,¹² J. Heefner,¹² Y. Hefetz,¹³ G. Heinzel,² I. S. Heng,³¹ M. Hennessy,²⁶ N. Hepler,²⁸ A. Heptonstall,³⁵ M. Heurs,³¹ M. Hewitson,² S. Hild,² N. Hindman,¹⁴ P. Hoang,¹² J. Hough,³⁵ M. Hrynevych,^{12,q} W. Hua,²⁶ M. Ito,³⁷ Y. Itoh,¹ A. Ivanov,¹² O. Jennrich,^{35,r} B. Johnson,¹⁴ W. W. Johnson,¹⁶ W. R. Johnston,²⁹ D. I. Jones,²⁸ L. Jones,²⁸ D. Jungwirth,^{12,s} V. Kalogera,²³ E. Katsavounidis,¹³ K. Kawabe,¹⁴ S. Kawamura,²² W. Kells,¹² J. Kern,^{15,t} A. Khan,¹⁵ S. Killbourn,³⁵ C. J. Killow,³⁵ C. Kim,²³ C. King,¹² P. King,¹² S. Klimentko,³⁴ S. Koranda,³⁹ K. Kötter,³¹ J. Kovalik,^{15,b} D. Kozak,¹² B. Krishnan,¹ M. Landry,¹⁴ J. Langdale,¹⁵ B. Lantz,²⁶ R. Lawrence,¹³ A. Lazzarini,¹² M. Lei,¹² I. Leonor,³⁷ K. Libbrecht,¹² A. Libson,⁸ P. Lindquist,¹² S. Liu,¹² J. Logan,^{12,u} M. Lormand,¹⁵ M. Lubinski,¹⁴ H. Lück,^{2,31} T. T. Lyons,^{12,u} B. Machenschalk,¹ M. MacInnis,¹³ M. Mageswaran,¹² K. Mailand,¹² W. Majid,^{12,b} M. Malec,^{2,31} F. Mann,¹² A. Marin,^{13,v} S. Márka,^{12,w} E. Maros,¹² J. Mason,^{12,x} K. Mason,¹³ O. Matherny,¹⁴ L. Matone,¹⁴ N. Mavalvala,¹³ R. McCarthy,¹⁴ D. E. McClelland,³ M. McHugh,¹⁸ J. W. C. McNabb,²⁸ G. Mendell,¹⁴ R. A. Mercer,³³ S. Meshkov,¹² E. Messaritaki,³⁹ C. Messenger,³³ V. P. Mitrofanov,²⁰ G. Mitselmakher,³⁴ R. Mittleman,¹³ O. Miyakawa,¹² S. Miyoki,^{12,y} S. Mohanty,²⁹ G. Moreno,¹⁴ K. Mossavi,² G. Mueller,³⁴ S. Mukherjee,²⁹ P. Murray,³⁵ J. Myers,¹⁴ S. Nagano,² T. Nash,¹² R. Nayak,¹¹ G. Newton,³⁵ F. Nocera,¹² J. S. Noel,⁴⁰ P. Nutzman,²³ T. Olson,²⁴ B. O'Reilly,¹⁵ D. J. Ottaway,¹³ A. Ottewill,^{39,z} D. Ouimette,^{12,s} H. Overmier,¹⁵ B. J. Owen,²⁸ Y. Pan,⁶ M. A. Papa,¹ V. Parameshwaraiah,¹⁴ A. Parameswaran,¹ C. Parameswariah,¹⁵ M. Pedraza,¹² S. Penn,¹⁰ M. Pitkin,³⁵ M. Plissi,³⁵ R. Prix,¹ V. Quetschke,³⁴ F. Raab,¹⁴ H. Radkins,¹⁴ R. Rahkola,³⁷ M. Rakhmanov,³⁴ S. R. Rao,¹² K. Rawlins,¹³ S. Ray-Majumder,³⁹ V. Re,³³ D. Redding,^{12,b} M. W. Regehr,^{12,b} T. Regimbau,⁷ S. Reid,³⁵ K. T. Reilly,¹² K. Reithmaier,¹² D. H. Reitze,³⁴ S. Richman,^{13,aa} R. Riesen,¹⁵ K. Riles,³⁶ B. Rivera,¹⁴ A. Rizzi,^{15,ab} D. I. Robertson,³⁵ N. A. Robertson,^{26,35} L. Robison,¹² S. Roddy,¹⁵ J. Rollins,¹³ J. D. Romano,⁷ J. Romie,¹² H. Rong,^{34,n} D. Rose,¹² E. Rothhoff,²⁸ S. Rowan,³⁵ A. Rüdiger,² P. Russell,¹² K. Ryan,¹⁴ I. Salzman,¹² V. Sandberg,¹⁴ G. H. Sanders,^{12,ac} V. Sannibale,¹² B. Sathyaprakash,⁷ P. R. Saulson,²⁷ R. Savage,¹⁴ A. Sazonov,³⁴ R. Schilling,² K. Schlaufman,²⁸ V. Schmidt,^{12,ad} R. Schnabel,¹⁹ R. Schofield,³⁷ B. F. Schutz,^{1,7} P. Schwinberg,¹⁴ S. M. Scott,³ S. E. Seader,⁴⁰ A. C. Searle,³ B. Sears,¹² S. Seel,¹² F. Seifert,¹⁹ A. S. Sengupta,¹¹ C. A. Shapiro,^{28,ae} P. Shawhan,¹² D. H. Shoemaker,¹³ Q. Z. Shu,^{34,af} A. Sibley,¹⁵ X. Siemens,³⁹ L. Sievers,^{12,b} D. Sigg,¹⁴ A. M. Sintes,^{1,32} J. R. Smith,² M. Smith,¹³ M. R. Smith,¹² P. H. Sneddon,³⁵ R. Spero,^{12,b} G. Stapfer,¹⁵ D. Steussy,⁸ K. A. Strain,³⁵ D. Strom,³⁷ A. Stuver,³⁸ T. Summerscales,²⁸ M. C. Sumner,¹² P. J. Sutton,¹² J. Sylvestre,^{12,ag} A. Takamori,¹² D. B. Tanner,³⁴ H. Tariq,¹² I. Taylor,⁷ R. Taylor,³⁵ R. Taylor,¹² K. A. Thorne,²⁸ K. S. Thorne,⁶ M. Tibbits,²⁸ S. Tilav,^{12,ah} M. Tinto,^{4,b} K. V. Tokmakov,²⁰ C. Torres,²⁹ C. Torrie,¹² G. Traylor,¹⁵ W. Tyler,¹² D. Ugolini,³⁰ C. Ungarelli,³³ M. Vallisneri,^{6,ai} M. van Putten,¹³ S. Vass,¹² A. Vecchio,³³ J. Veitch,³⁵ C. Vorvick,¹⁴ S. P. Vyachanin,²⁰ L. Wallace,¹² H. Walther,¹⁹ H. Ward,³⁵ B. Ware,^{12,b} K. Watts,¹⁵ D. Webber,¹² A. Weidner,¹⁹ U. Weiland,³¹

A. Weinstein,¹² R. Weiss,¹³ H. Welling,³¹ L. Wen,¹² S. Wen,¹⁶ J. T. Whelan,¹⁸ S. E. Whitcomb,¹² B. F. Whiting,³⁴
 S. Wiley,⁵ C. Wilkinson,¹⁴ P. A. Willems,¹² P. R. Williams,^{1,aj} R. Williams,⁴ B. Willke,³¹ A. Wilson,¹²
 B. J. Winjum,^{28,e} W. Winkler,² S. Wise,³⁴ A. G. Wiseman,³⁹ G. Woan,³⁵ D. Woods,¹ R. Wooley,¹⁵ J. Worden,¹⁴
 W. Wu,³⁴ I. Yakushin,¹⁵ H. Yamamoto,¹² S. Yoshida,²⁵ K. D. Zaleski,²⁸ M. Zanolin,¹³ I. Zawischa,^{31,ak} L. Zhang,¹¹
 R. Zhu,¹ N. Zotov,¹⁷ M. Zucker,¹⁵ and J. Zweizig¹²

(LIGO Scientific Collaboration)

¹*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476 Golm, Germany*

²*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany*

³*Australian National University, Canberra, 0200, Australia*

⁴*California Institute of Technology, Pasadena, California 91125, USA*

⁵*California State University Dominguez Hills, Carson, California 90747, USA*

⁶*Caltech-CaRT, Pasadena, California 91125, USA*

⁷*Cardiff University, Cardiff, CF2 3YB, United Kingdom*

⁸*Carleton College, Northfield, Minnesota 55057, USA*

⁹*Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*

¹⁰*Hobart and William Smith Colleges, Geneva, New York 14456, USA*

¹¹*Inter-University Centre for Astronomy and Astrophysics, Pune-411007, India*

¹²*LIGO—California Institute of Technology, Pasadena, California 91125, USA*

¹³*LIGO—Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

¹⁴*LIGO Hanford Observatory, Richland, Washington 99352, USA*

¹⁵*LIGO Livingston Observatory, Livingston, Louisiana 70754, USA*

¹⁶*Louisiana State University, Baton Rouge, Louisiana 70803, USA*

¹⁷*Louisiana Tech University, Ruston, Louisiana 71272, USA*

¹⁸*Loyola University, New Orleans, Louisiana 70118, USA*

¹⁹*Max Planck Institut für Quantenoptik, D-85748, Garching, Germany*

²⁰*Moscow State University, Moscow, 119992, Russia*

²¹*NASA/Goddard Space Flight Center, Greenbelt, Maryland 20771, USA*

²²*National Astronomical Observatory of Japan, Tokyo 181-8588, Japan*

²³*Northwestern University, Evanston, Illinois 60208, USA*

²⁴*Salish Kootenai College, Pablo, Montana 59855, USA*

²⁵*Southeastern Louisiana University, Hammond, Louisiana 70402, USA*

²⁶*Stanford University, Stanford, California 94305, USA*

²⁷*Syracuse University, Syracuse, New York 13244, USA*

²⁸*The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

²⁹*The University of Texas at Brownsville and Texas Southmost College, Brownsville, Texas 78520, USA*

³⁰*Trinity University, San Antonio, Texas 78212, USA*

³¹*Universität Hannover, D-30167 Hannover, Germany*

³²*Universitat de les Illes Balears, E-07122 Palma de Mallorca, Spain*

³³*University of Birmingham, Birmingham, B15 2TT, United Kingdom*

³⁴*University of Florida, Gainesville, Florida 32611, USA*

³⁵*University of Glasgow, Glasgow, G12 8QQ, United Kingdom*

³⁶*University of Michigan, Ann Arbor, Michigan 48109, USA*

³⁷*University of Oregon, Eugene, Oregon 97403, USA*

³⁸*University of Rochester, Rochester, New York 14627, USA*

³⁹*University of Wisconsin—Milwaukee, Milwaukee, Wisconsin 53201, USA*

⁴⁰*Washington State University, Pullman, Washington 99164, USA*

(Received 30 May 2005; published 25 October 2005)

We use 373 hours (≈ 15 days) of data from the second science run of the LIGO gravitational-wave detectors to search for signals from binary neutron star coalescences within a maximum distance of about 1.5 Mpc, a volume of space which includes the Andromeda Galaxy and other galaxies of the Local Group of galaxies. This analysis requires a signal to be found in data from detectors at the two LIGO sites, according to a set of coincidence criteria. The background (accidental coincidence rate) is determined from the data and is used to judge the significance of event candidates. No inspiral gravitational-wave events were identified in our search. Using a population model which includes the Local Group, we establish an upper limit of less than 47 inspiral events per year per Milky Way equivalent

galaxy with 90% confidence for nonspinning binary neutron star systems with component masses between 1 and $3M_{\odot}$.

DOI: [10.1103/PhysRevD.72.082001](https://doi.org/10.1103/PhysRevD.72.082001)

PACS numbers: 95.85.Sz, 04.80.Nn, 07.05.Kf, 97.80.-d

I. INTRODUCTION

The search for gravitational waves has entered a new era with the scientific operation of kilometer scale laser interferometers. These L-shaped instruments are sensitive to minute changes in the relative lengths of their orthogonal arms that would be produced by gravitational waves [1]. The Laser Interferometer Gravitational-wave Observatory (LIGO) [2,3] consists of three Fabry-Perot-Michelson interferometers: two interferometers are housed at the site in Hanford, WA; a single interferometer is housed in Livingston, Louisiana. In 2003, all three instruments simultaneously collected data under stable operating conditions during two science runs. Even though the instruments

were not yet performing at their design sensitivity, the data represent the best broadband sensitivity to gravitational waves that has been achieved to date.

In this paper, we report the methods and results of a search for gravitational waves from binary neutron star systems, using data from the science run conducted early in 2003. These waves are expected to be emitted at frequencies detectable by LIGO during the final few seconds of inspiral as the binary orbit decays due to the loss of energy in gravitational radiation [4]. A previous search [5], using data from the first LIGO science run, reported an upper limit on the rate of coalescences within our Galaxy and the magellanic clouds. This paper uses an analysis pipeline which is optimized for detection by using data only from times when interferometers were operating properly at both LIGO sites. By demanding that a gravitational wave be seen at both sites, we strongly suppress the rate of background events from nonastrophysical disturbances. Moreover, this approach allows us to judge the significance of any apparent event candidate in the context of the background distribution, which is also determined from the data.

The search described here has essentially perfect efficiency for detecting binary neutron star inspirals within the Milky Way and the magellanic clouds (as measured by Monte-Carlo simulations), and could detect some inspirals as far away as the Andromeda and Triangulum galaxies (M31 and M33). The rate of coalescences in these galaxies, based on the population of known binary neutron star systems [6], is expected to be very low, so that a detection by the present search would be highly surprising. In fact, no coincident event candidates were observed in excess of the measured background. The data are therefore used to place an improved direct observational upper limit on the rate of binary neutron star coalescence events in the Universe.

II. DATA SAMPLE

The LIGO Hanford Observatory (LHO) in Washington state has two independent detectors sharing a common vacuum envelope, one with 4 km long arms (H1) and one with 2 km long arms (H2). The LIGO Livingston Observatory (LLO) in Louisiana has one detector with 4 km long arms (L1). All three detectors operated during the second LIGO science run, referred to as S2, which spanned 59 days from February 14 to April 14, 2003. During operation, feedback to the mirror positions and to the laser frequency keeps the optical cavities near resonance, so that interference in the light from the two arms recombining at the beam splitter is strongly dependent on

^aCurrently at Stanford Linear Accelerator Center.

^bCurrently at Jet Propulsion Laboratory.

^cPermanent Address: HP Laboratories.

^dCurrently at Rutherford Appleton Laboratory.

^eCurrently at University of CA, Los Angeles.

^fCurrently at Hofstra University.

^gPermanent Address: GReCO, Institut d'Astrophysique de Paris (CNRS).

^hCurrently at La Trobe University, Bundoora VIC, Australia.

ⁱCurrently at Keck Graduate Institute.

^jCurrently at National Science Foundation.

^kCurrently at University of Sheffield.

^lCurrently at Ball Aerospace Corporation.

^mCurrently at European Gravitational Observatory.

ⁿCurrently at Intel Corp.

^oCurrently at University of Tours, France.

^pCurrently at Lightconnect Inc.

^qCurrently at W. M. Keck Observatory.

^rCurrently at ESA Science and Technology Center.

^sCurrently at Raytheon Corporation.

^tCurrently at NM Institute of Mining and Technology/Magdalena Ridge Observatory Interferometer.

^uCurrently at Mission Research Corporation.

^vCurrently at Harvard University.

^wPermanent Address: Columbia University.

^xCurrently at Lockheed-Martin Corporation.

^yPermanent Address: University of Tokyo, Institute for Cosmic Ray Research.

^zPermanent Address: University College Dublin.

^{aa}Currently at Research Electro-Optics Inc.

^{ab}Currently at Institute of Advanced Physics, Baton Rouge, LA.

^{ac}Currently at Thirty Meter Telescope Project at Caltech.

^{ad}Currently at European Commission, DG Research, Brussels, Belgium.

^{ae}Currently at University of Chicago.

^{af}Currently at LightBit Corporation.

^{ag}Permanent Address: IBM Canada Ltd.

^{ah}Currently at University of DE.

^{ai}Permanent Address: Jet Propulsion Laboratory.

^{aj}Currently at Shanghai Astronomical Observatory.

^{ak}Currently at Laser Zentrum Hannover.

the difference between the lengths of the two arms. A photodiode at the antisymmetric port of the detector senses this light, and a digitized signal is recorded at a sampling rate of 16 384 Hz. This channel can then be searched for a gravitational-wave signal. More details on the detectors' instrumental configuration and performance can be found in [3,7].

While the detailed noise spectrum of a detector affects different gravitational-wave searches in different ways, we can summarize the sensitivity of a detector for low-mass inspiral signals in terms of the range for an archetypal source. Specifically, the range is the distance at which an optimally oriented and located binary system¹ with the mass of each component equal to $1.4M_{\odot}$ would yield an amplitude signal-to-noise ratio (SNR) of 8 when extracted from the data using optimal filtering. During the first LIGO science run (referred to as S1), the L1 detector had the greatest range, typically 0.18 Mpc. During the S2 run, all three detectors were substantially more sensitive than this, with ranges of 2.0, 0.9, and 0.6 Mpc for L1, H1, and H2 averaged over all times during the run. Typical amplitude spectral densities of detector noise are shown in Fig. 1.

The amount of science data with good performance and stable operating conditions was limited by environmental factors (especially high ground motion at LLO and strong winds at LHO), occasional equipment failures, and periodic special investigations. Over the 1415 h duration of the S2 run, the total amount of science data obtained was 536 hours for L1, 1044 hours for H1, and 822 hours for H2.

The analysis presented here uses data collected while the LLO detector was operating at the same time as one or both of the LHO detectors. Science data during which both H1 and H2 were operating but L1 was not, amounting to 385 hours, was not used in this analysis because of concerns about possible environmentally induced correlations between the data streams of these two colocated detectors; this data set, as well as data collected while only one of the LIGO detectors was in science mode, will be used in a separate analysis together with data from the TAMA300 detector [8], which conducted "Data Taking 8" concurrently with the LIGO S2 run.

Of all the data, approximately 9% (uniformly sampled within the run so as to be representative of the whole data set) was used as *playground* data for tuning parameters and thresholds in the analysis pipeline, and for use in identifying vetoes that were effective in eliminating spurious events. This playground data set was excluded from the gravitational-wave inspiral upper limit calculation because event selection and pipeline tuning, as described in Secs. IV, V, and VI, introduces statistical bias which cannot be accounted for. The playground data set was searched for




FIG. 1 (color online). Typical sensitivities, expressed as amplitude spectral densities of detector noise converted to equivalent gravitational-wave strain, for the best detector during the S1 run (the Livingston detector) and for all three detectors during the S2 run. The solid lower line is the design sensitivity for the LIGO 4-km detectors; the dashed line is the design sensitivity for the LIGO 2-km detector at Hanford.

inspiral signals, however, so a potential detection during these times was not excluded. After applying data quality cuts (as detailed in Sec. V) and accounting for short time intervals which could not be searched for inspiral signals by the filtering algorithm (described in Sec. IV), the observation time consisted of 242 hours of triple-detector data, plus 99 hours of L1-H1 and 32 hours of L1-H2 data, for a total observation time of 373 hours. For the upper limit result, the nonplayground observation time was 339 hours. A summary of the amount of single-, double-, and triple-detector times is provided in Fig. 2.




FIG. 2. The number of hours that each detector combination was operational during the S2 run. The upper number gives the amount of time the specific instruments were coincidentally operational. The lower number gives the total time that was searched for inspiral triggers. The shaded region corresponds to the data used in this search.

¹An optimally oriented and located binary system would be located at the detector's zenith with its orbital plane perpendicular to the line of sight between the detector and the source.

As with earlier analyses of LIGO data [5], the output of the antisymmetric port of the detector was calibrated to obtain a measure of the relative strain $s = \Delta L/L$ of the detector arms, where $\Delta L = L_x - L_y$ is the difference in length between one arm (the x arm) and the other (the y arm), and L is the average arm length. Reference calibration functions, tracing out the frequency-dependent response of the detectors, were measured (by moving the end mirrors of the detector with a known displacement) before and after the science run, and once during the science run; all three measurements gave consistent results. The changing optical gain of the calibration was monitored continuously during the run by applying sinusoidal motions with fixed frequency to the end mirrors. This continuous monitoring, averaged once a minute, allowed for small corrections to the calibration due to loss of light power in the arms, which can be caused by drifting optical alignment.

III. TARGET SOURCES

Binary neutron star (BNS) systems in the Milky Way, known from radio pulsar observations [9], provide indirect evidence for the existence of gravitational waves [10]. Based on current astrophysical understanding, the spatial distribution of binary neutron stars is expected to follow that of star formation in the Universe. A measure of star formation is the blue luminosity of galaxies appropriately corrected for dust extinction and reddening. Therefore we model the spatial distribution of double neutron stars according to the corrected blue-light distribution of nearby galaxies [11]. While the masses of neutron stars in the few known binary systems are all near $1.4M_\odot$, population synthesis simulations suggest that some systems will have component masses as low as $\sim 1M_\odot$ and as high as the theoretical maximum neutron star mass of $\sim 3M_\odot$ [12]. Thus, we search for inspiral signals from binary systems with component masses in this range. Note that the higher-mass systems radiate more energy in gravitational waves and can thus be detected at a greater distance at a given SNR. For component masses below $1M_\odot$, a search is reported in Ref. [13].

When the LIGO detectors reach their design sensitivities, they will be capable of detecting inspiral signals from thousands of galaxies, reaching beyond the Virgo Cluster for systems with optimal location and orientation. At that sensitivity, the rate of detectable binary neutron star coalescences could be as high as 0.7 per year, though it is more likely to be an order of magnitude smaller [6]. For this analysis, our target population includes the Milky Way and all significant galaxies within a distance of 3 Mpc, which is roughly the maximum distance for which a $3\text{--}3M_\odot$ inspiral could be detected in coincidence by the L1 and H1 interferometers with a SNR of 6 in H1. This population includes the Local Group of galaxies, whose total blue luminosity is dominated by the Andromeda

Galaxy (M31), as well as some galaxies from neighboring groups. We cannot hope to detect *all* inspirals within this volume, because most systems in the population have lower masses and because the received signal amplitude is reduced, on average, depending on the orientation and location of the source relative to the detector.

Table I gives the parameters we use for the galaxies in the target population out to 1.5 Mpc, i.e., the maximum distance for which we had a nonzero detection efficiency in our simulations. The coordinates and distances of the galaxies in Table I are taken from a catalog by Mateo [14], when available; this catalog is favored because distances quoted are from individual, focused studies of each of the nearby galaxies he includes. The rest of the distances, with only 100 kpc accuracy, are taken from the Tully Nearby Galaxies catalog [15]. Data for blue luminosities are derived from the apparent blue magnitudes (corrected for reddening) quoted in Ref. [16], and the distances shown in Table I. We measured the efficiency of our search using Monte-Carlo simulations, where the sources in the target population had a mass distribution as described in Ref. [12], following the same guidelines as in the population models used in Ref. [17]. We used simulations with a population of neutron stars from galaxies up to 3 Mpc away, overextending our target population, although we did not detect any simulated injections from sources farther than 1.5 Mpc away.

Within the LIGO frequency band, the gravitational waveform produced by binary neutron star systems is well described by the restricted second-post-Newtonian approximation [18–20]. The spins of the neutron stars are not expected to significantly affect the orbital motion or the waveform [18]. Tidal coupling and other finite-body effects dependent on the equation of state also are not expected to significantly affect the waveform in the LIGO band [21]. The waveform received at Earth is therefore parametrized by the masses of the companions, the distance to the binary, the inclination of the system to the plane of the sky, and by the initial orbital phase when the waveform enters the LIGO band. The waveforms consist of two polarizations, the plus (h_+) and cross (h_\times) polarizations, which describe the two orthogonal tidal distortions produced by the waves. These polarization basis states are defined with respect to the orientation of the binary orbit relative to the line of sight. An interferometric detector is sensitive to a particular linear combination of these two polarizations; this is described by two response functions F_+ and F_\times so that the expected gravitational-wave signal is

$$h(t) = F_+ h_+(t) + F_\times h_\times(t). \quad (1)$$

The response functions depend on the location of the detector on the Earth and on the orientation of the detector's arms. The interferometers at LHO and LLO are

TABLE I. The galaxies in our population within 1.5 Mpc. The next to last column indicates the number of injections detected in coincidence (with signal-to-noise-ratio defined in Eq. (10) satisfying $\rho^2 > 89$), over the number of injections performed in our simulations. The last column indicates the cumulative number of equivalent Milky Way galaxies contributed by systems within the corresponding distance, calculated from the search efficiency and the blue-light luminosity.

Name	Right ascension (Hour:Min)	Declination (Deg:Min)	Distance (kpc)	Blue-light luminosity relative to Milky Way	Detected/injected with $\rho^2 > 89$	Cumulative N_G
Milky Way	1	686/686	1
LMC	+05:23.6	-69:45	49	0.128	57/57	1.128
SMC	+00:52.7	-72:50	58	0.037	8/8	1.165
NGC6822	+19:44.9	-14:49	490	0.01	0/4	1.165
NGC185	+00:39.0	+48:20	620	0.007	1/3	1.1673
M110	+00:40.3	+41:41	815	0.018	5/100	1.1682
M31	+00:42.7	+41:16	770	2.421	108/1791	1.3142
M32	+00:42.7	+40:52	805	0.019	9/129	1.3155
IC10	+00:20.4	+59:18	825	0.031	1/10	1.3186
M33	+01:33.9	+30:39	840	0.319	9/219	1.3318
NGC300	+00:54.9	-37:41	1200	0.052	1/40	1.3331
M81	+09:55.6	+69:04	1400	0.196	1/147	1.3344
NGC55	+00:14.9	-39:11	1480	0.175	2/122	1.3373

aligned as closely as possible, but the curvature of the Earth causes a slight difference in their antenna patterns.

IV. FILTERING AND TRIGGER GENERATION

We generated event *triggers* by filtering the data $s(t)$ from each detector with matched filters designed to detect the expected signals. For any given binary neutron star mass pair, $\{m_1, m_2\}$, we constructed the expected frequency-domain inspiral waveform template, $\tilde{h}(f)$, using a stationary-phase approximation to the restricted second-post-Newtonian waveform [22].² Here, the tilde indicates the Fourier transform of a time series $h(t)$ according to the convention

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

The matched filter output is then the complex time series

$$z(t) = x(t) + iy(t) = 4 \int_0^{\infty} \frac{\tilde{h}^*(f) \tilde{s}(f)}{S(f)} e^{2\pi i f t} df \quad (3)$$

where $x(t)$ is the (real) matched filter output for the inspiral waveform with a zero orbital phase, and $y(t)$ is the (real) matched filter output for the inspiral waveform with a $\pi/2$ orbital phase. The quantity $S(f)$ is the one-sided strain noise power spectral density, estimated from the data. The matched filter variance is given by

$$\sigma^2 = 4 \int_0^{\infty} \frac{|\tilde{h}(f)|^2}{S(f)} df \quad (4)$$

²The stationary-phase approximation to the Fourier transform of inspiral template waveforms was shown to be sufficiently accurate for gravitational-wave detection in Ref. [23].

which depends on the template's amplitude normalization. The amplitude SNR is then

$$\rho = |z|/\sigma. \quad (5)$$

The templates were normalized to a binary neutron star inspiral at an *effective distance* of 1 Mpc, where the effective distance of a waveform is the distance for which a binary neutron star system would produce the waveform if it were optimally oriented. Thus,

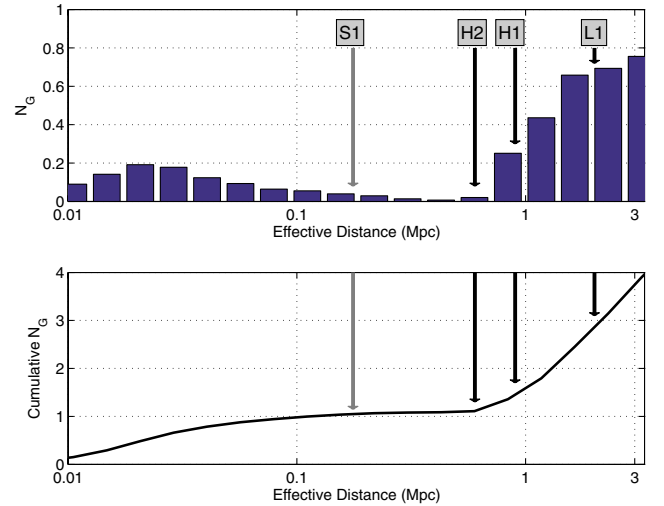


FIG. 3 (color online). The upper panel shows the histogram of the number of Milky Way equivalent galaxies (N_G) as a function of *effective distance*. Labeled arrows indicate the average range of the best detector in the first science run (S1), as well as the average range of each detector during the second science run. The cumulative total within a given effective distance is shown in the bottom panel.

$$D_{\text{eff}} = \frac{\sigma}{\rho} \quad (6)$$

is an estimate of the effective distance, in Mpc, of a putative signal that produces SNR ρ . The number of Milky Way equivalent galaxies as a function of effective distance is shown in Fig. 3.

Since each binary neutron star mass pair $\{m_1, m_2\}$ would produce a slightly different waveform, we constructed a *bank* of templates with different mass pairs such that, for any actual mass pair with $1M_{\odot} \leq m_2 \leq m_1 \leq 3M_{\odot}$, the loss of SNR due to the mismatch of the true waveform from that of the best fitting waveform in the bank is less than 3% [24,25].

Although a threshold on the matched filter output ρ would be the optimal detection criterion for an accurately known inspiral waveform in the case of stationary, Gaussian noise, the character of the data used in this analysis is known to be neither stationary nor Gaussian. Indeed, many classes of transient instrumental artifacts have been categorized, some of which produce copious numbers of spurious, large SNR events. In order to reduce the number of spurious event triggers, we adopted a now-standard χ^2 requirement [26]: The matched filter template is divided into p frequency bands, which are chosen so that each band would contribute a fraction $1/p$ of the total SNR if a true signal (and no detector noise) were present. In our analysis, we used $p = 15$, as explained in Sec. VIB. We then construct a chi-squared statistic comparing the magnitude and phase of SNR accumulated in each band to the expected amount. For a true signal in Gaussian noise, the resulting statistic is χ^2 distributed with $2p - 2$ degrees of freedom (since the SNR is constructed out of the two matched filters x and y which are both measured). Instrumental artifacts tend to produce very large χ^2 values and can be rejected by requiring χ^2 to be less than some reasonable threshold. Because of the discreteness of the template bank, however, a real signal will generally not match precisely the nearest template in the bank. Consequently, χ^2 has a noncentral chi-squared distribution with a noncentrality parameter $\lambda \leq (p - 1)\rho^2\mu^2$, where μ is the fractional loss of SNR due to mismatch between template and signal [26]. For sufficiently small μ and moderate values of ρ , this effect is not important, but when ρ is large, it must be taken into account. This is done by applying a threshold with the following parametrization:

$$\chi^2 \leq (p + \delta^2\rho^2)\xi^*. \quad (7)$$

The threshold multiplier ξ^* , the number of bins p , and the value of δ were determined by tuning on the playground data, as will be described in Sec. VIB.

In this analysis, the SNR $\rho(t)$ was computed for each template in the bank [27]. Whenever $\rho(t)$ exceeded a threshold ρ^* , the value of χ^2 was computed for that time. If χ^2 was below the threshold in Eq. (7), then the local maximum of $\rho(t)$ was recorded as a *trigger*. Each trigger is

represented by a vector of values: the masses which define the template, the maximum value of ρ and corresponding value of χ^2 , the inferred coalescence time, the effective distance D_{eff} , and the coalescence phase $\phi_0 = \tan^{-1}(y/x)$.

V. DATA QUALITY CHECKS AND VETOES

In practice, the performance of the matched filtering algorithm described above is influenced by nonstationary optical alignment, servo control settings, and environmental conditions. We used two strategies to avoid problematic data in this search. One was to evaluate *data quality* over relatively long time intervals, using several different tests. Time intervals identified as being suspect or demonstrably bad were skipped when filtering the data. The other method was to look for signatures in environmental monitoring channels and auxiliary interferometer channels which would indicate an external disturbance or instrumental glitch (a large transient fluctuation), allowing us to *veto* any triggers recorded at that time.

Several data quality tests were applied *a priori*, leading us to omit data when calibration information was missing or unreliable, servo control settings were not at their nominal values, or there were input/output controller timing problems. Additional tests were performed to characterize the noise level in the interferometer in various frequency bands and to check for problems with the photodiodes and associated electronics. The playground data set was used to judge the relevance of these additional tests, and two data quality tests were found to correlate with inspiral triggers found in the playground. One of these pertained only to the H1 interferometer; there were occasional, abnormally high noise levels in the H1 antisymmetric port channel that were apparent when this signal was averaged over a minute. Data was rejected only if this excessive noise was present for at least 3 consecutive minutes. The other data quality test used to reject data pertained to saturation of the photodiode at the antisymmetric port. These photodiode saturation events correlated with a small, but significant, number of L1 inspiral triggers. We required the absence of photodiode saturation in the data from all three detectors.

A gravitational wave would be most evident in the signal obtained at the antisymmetric port. We examined many other auxiliary interferometer channels, which monitor the light in the interferometer at points other than the antisymmetric port, to look for correlations between glitches found in the readouts of these channels and inspiral event triggers found in the playground data. These auxiliary channels are sensitive to certain instrumental artifacts which may also affect the antisymmetric port, and therefore may provide highly effective veto conditions. Although these channels are expected to have little or no sensitivity to a gravitational wave, we considered the possibility that an actual astrophysical signal could produce the observed glitches in the auxiliary channel due to some physical or electronic

coupling. This possibility was tested by means of *hardware* injections, in which a simulated inspiral signal is injected into the data by physically moving one of the end mirrors of the interferometer. These hardware injections were used for validation of the analysis pipeline, as described in Sec. VIC. Unlike the software injections that are used to measure the pipeline efficiency (described in Sec. VIB), these hardware injections allow us to establish a limit on the effect that a true signal would have on the auxiliary channels. Only those channels that were unaffected by the hardware injections were considered safe for use as potential veto channels.

We used an analysis program, *glitchMon*, [28] to identify large amplitude transient signals in auxiliary channels. Numerous channels were examined by *glitchMon*, which generates a list of times when glitches occurred, identified by a filtered time series crossing a chosen threshold. A veto condition based on a given list of *glitchMon* triggers was defined by choosing a fixed time window around each glitch and rejecting any inspiral event trigger with a coalescence time within the window.

For each veto condition considered, we evaluated the veto efficiency (percentage of inspiral events eliminated), use percentage (percentage of veto triggers which veto at least one inspiral event), and dead time (percentage of science-data time eliminated by the veto). Once a channel was identified, some tuning was done of the filters used, and the thresholds and time windows chosen, to optimize the efficiency, especially for high SNR candidates, without an excessive dead time. The parameter tuning was done only with the inspiral triggers found in the playground.

No efficient candidate veto channels were identified for H1 and H2; there were some candidates for L1. Nonstationary noise in the low-frequency part of the sensitivity range used for the inspiral search appeared to be the dominant cause for glitch events in the data. The original frequency range for the binary neutron star inspiral search extended from 50 Hz to 2048 Hz. It was discovered that many of the L1 inspiral triggers appeared to be the result of nonstationary noise with frequency content around 70 Hz. An important auxiliary channel, L1:LSC-POB_I, proportional to the length fluctuations of the power recycling cavity, was found to have highly variable noise at 70 Hz. As a consequence, it was decided that the low-frequency cutoff of the binary neutron star inspiral search should be increased from 50 Hz to 100 Hz. This subsequently reduced the number of inspiral triggers. An inspection of artificial signals injected into the data revealed a very small loss of efficiency for binary neutron star inspiral signal detection resulting from the increase in the low-frequency cutoff.

Even after raising the low-frequency cutoff, the L1:LSC-POB_I channel was found to be an effective veto when filtered appropriately. Because of the characteristics of the inspiral templates' response to large glitches, we decided to

veto with a very wide time window, -4 s to $+8$ s, around the time of each L1:LSC-POB_I trigger. With this choice, 12% of the BNS inspiral triggers with $\rho > 8$ in the playground were vetoed, as well as 5 out of the 9 triggers found in the playground with $\rho > 10$. The use percentage of the veto triggers was 18% for $\rho > 8$ and 0.7% for $\rho > 10$, whereas values of 3% and $<0.1\%$ (respectively) would be expected from random coincidences if the veto triggers had no real correlation with the inspiral triggers. The total L1 dead time using this veto condition in the playground was 2.7%. The performance in the full data set was consistent with that found in the playground: the final observation time (including playground) was reduced by this veto from 385 hours to 373 hours. A more extensive discussion of LIGO's S2 binary inspiral veto study can be found in [29].

VI. SEARCH FOR COINCIDENT EVENT CANDIDATES

A. Analysis pipeline

The detection of a gravitational-wave inspiral signal in the S2 data would (at the least) require triggers in L1 and one or more of the Hanford instruments with consistent arrival times (separated by the light travel time between the detectors) and waveform parameters. Requiring temporal coincidence between the two observatories greatly reduces the background rate due to spurious triggers, thus allowing an increased confidence for detection candidates. When detectors at both observatories are operating simultaneously, we may obtain an estimate of the rate of background triggers by time shifting the Hanford triggers with respect to the Livingston triggers and applying the same coincidence requirements to the time-shifted triggers, as described in Sec. VII. In this way, we can measure the rate of accidental coincidences in our search.

During the S2 run, the three LIGO detectors had substantially different sensitivities. The sensitivity of the L1 detector was greater than that of either Hanford detector throughout the run. Since the orientations of the LIGO interferometers are similar, we expect that signals of astrophysical origin detected in the Hanford interferometers generally are detectable in the L1 interferometer. Using this as a guiding principle, we have constructed a triggered search pipeline, summarized in Fig. 4. We search for inspiral triggers in the most sensitive interferometer (L1), and only when a trigger is found in this interferometer do we search for a coincident trigger in the less sensitive interferometers. This approach reduces significantly the computational power necessary to perform the search, without compromising the detection efficiency of the pipeline.

Times when each interferometer was in stable operation are called science segments. The science segments from each interferometer are used to construct three data sets corresponding to: (1) times when all three interferometers were operating, (2) times when only the L1 and H1 inter-

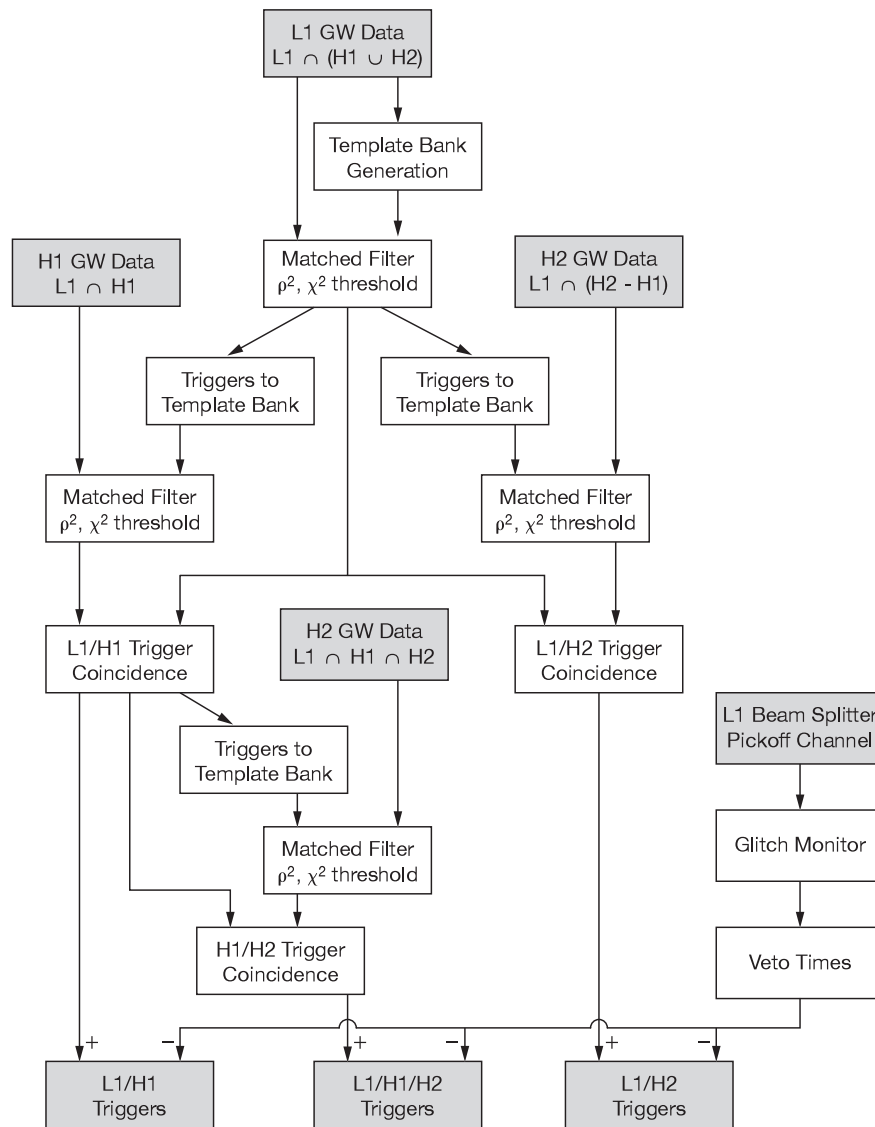


FIG. 4. The inspiral analysis pipeline used to determine the reported upper limit. $L1 \cap (H1 \cup H2)$ indicates times when the L1 interferometer was operating in coincidence with one or both of the Hanford interferometers. $L1 \cap H1$ indicates times when the L1 interferometer was operating in coincidence with the H1 interferometer. $L1 \cap (H2 - H1)$ indicates times when the L1 interferometer was operating in coincidence with only the H2 interferometer. The outputs of the search pipeline are triggers that belong to one of the two double-detector coincident data sets or to the triple-detector data set.

ferometers were operating, and (3) times when only the L1 and H2 interferometers were operating. The pipeline produces a list of coincident triggers for each of these three data sets as described below.

The science segments were analyzed in blocks of 2048 seconds using the FINDCHIRP implementation of matched filtering for inspiral signals in the LIGO Algorithm Library [30]. In this code, the data set for each 2048 s block is first down-sampled from 16384 Hz to 4096 Hz. It is subsequently high-pass filtered and a low-frequency cutoff of 100 Hz imposed. The calibrated instrumental response for the block is calculated using the average of the calibrations (measured every minute) over the duration of the block.

Triggers are not searched for within the first and last 64 s of a given block, so subsequent blocks are overlapped by 128 s to ensure that all of the data in a continuous science segment (except for the first and last 64 seconds) are searched for triggers. Any science segments shorter than 2048 s are ignored. If a science segment cannot be exactly divided into overlapping blocks (as is usually the case) the remainder of the data set is covered by a special 2048-s block which overlaps with the previous block as much as necessary to allow it to reach the end of the segment. For this final block, a parameter is set to restrict the inspiral search to the time interval not covered by any previous block, as shown in Fig. 5.




FIG. 5. The algorithm used to divide science segments into data analysis segments. Science segments are divided into 2048 s blocks overlapped by 128 s. (Science segments shorter than 2048 s are ignored.) An additional block with a larger overlap is added to cover any remaining data at the end of a science segment. Each block is divided into 15 data segments of length 256 s for filtering. The first and last 64 s of each data segment are ignored, so the segments overlap by 128 s. Areas shaded black are searched for triggers by the search pipeline. The gray area in the last block of the science segment is not searched for triggers as this time is covered by the preceding block, although these data points are used in estimating the noise power spectral density for the final block.

Each block is further split into 15 data segments of length 256 seconds overlapped by 128 seconds. The power spectrum $S(f)$ for the 2048 seconds of data is estimated by taking the median of the power spectra of the 15 segments. (We use the median and not the mean to avoid biased estimates due to large outliers, produced by nonstationary data.) The average calibration is applied to the data in each data segment, and the matched filter output in Eq. (3) is computed for each template in the template bank.

In order to avoid end effects when applying the matched filter, the frequency-weighting factor, nominally $1/S(f)$, is altered so that its inverse Fourier transform has a maximum duration of ± 16 seconds. The output of the matched filter near the beginning and end of each segment is corrupted by end effects due to the finite duration of the power spectrum weighting and also the inspiral template. By ignoring the filter output within 64 s of the beginning and end of each segment, we ensure that only uncorrupted filter output is searched for inspiral triggers. This necessitates the overlapping of segments and blocks as described above.

The power spectral density (PSD) of the noise in the Livingston detector is estimated independently for each L1 block that is coincident with operation of a Hanford detector [denoted $L1 \cap (H1 \cup H2)$]. The PSD is used to lay out a template bank for filtering that block, according to the parameters for mass ranges and minimal match [24,25]. The data from the L1 interferometer for the block are then filtered, using that bank, with a signal-to-noise threshold ρ_L^* and χ^2 veto threshold ξ_L^* to produce a list of triggers as described in Sec. IV. For each block in the Hanford inter-

ferometers, a *triggered bank* is created consisting of every template which produced at least one trigger in L1 during the time of the Hanford block. This is used to filter the data from the Hanford interferometers with signal-to-noise and χ^2 thresholds specific to the interferometer, giving a total of six thresholds that may be tuned. For times when only the H2 interferometer is operating in coincidence with L1 [denoted $L1 \cap (H2 - H1)$] the triggered bank is used to filter the H2 blocks that overlap with L1 data; these triggers are used to test for L1/H2 coincidence. All H1 data that overlaps with L1 data (denoted $L1 \cap H1$) are filtered using the triggered bank for that block. For H1 triggers produced during times when all three interferometers were operating, a second triggered bank is produced for each H2 block consisting of every template which produced at least one trigger found in coincidence in L1 and H1 during the time of the H2 block. The H2 block is filtered with this bank. Any H2 triggers found with this bank are tested for triple coincidence with the L1 and H1 triggers. The remaining triggers from H1, when H2 is not available, are used to search for L1/H1 coincident triggers.

For a trigger to be considered coincident between two interferometers, the following conditions must be fulfilled: (1) Triggers must be observed in both interferometers within a temporal coincidence window that allows for the error in measurement of the time of the trigger, δt . If the detectors are not colocated, this parameter is increased by the light travel time between the observatories (10 ms for traveling 3000 km at the speed of light). (2) We then ensure that the triggers have consistent waveform parameters by demanding that the two mass parameters for the template are identical to within an error of δm . (3) For H1 and H2, we may impose an amplitude cut on the triggers given by

$$\frac{|D_1 - D_2|}{D_1} < \frac{\epsilon}{\rho_2} + \kappa, \quad (8)$$

where D_1 (D_2) is the effective distance of the trigger in the first (second) detector and ρ_2 is the signal-to-noise ratio of the trigger in the second detector. The parameters ϵ and κ are tunable. As shown in Fig. 6, the nonperfect alignment of LLO and LHO (due to their different latitudes) can occasionally cause large variations in the detected signal amplitudes for astrophysical signals. In order to disable the amplitude cut when comparing triggers from LLO and LHO, we set $\kappa = 1000$.

If the detectors are at the same site, we ask if the maximum distance to which H2 can see at the signal-to-noise threshold ρ_{H2}^* is greater than the distance of the H1 trigger, allowing for errors in the measurement of the trigger distance. If this is the case, we demand time, mass, and effective distance coincidence. If the distance to which H2 can see overlaps the error in measured distance of the H1 trigger, we search for a trigger in H2, but always keep the H1 trigger even if no coincident trigger is found. If the minimum of the error in measured distance of




FIG. 6 (color online). Ratio of effective distance at the Hanford and Livingston Observatories for the injections from sources in Table I, versus Greenwich Mean Sidereal Time (GMST). The sharp feature near 21.5 hrs is due to M31 (Andromeda) passing through a sky position at the Hanford node, producing much larger effective distances at LHO than at LLO. The softer feature near 14 hrs is due to M33 (Triangulum Galaxy) passing through a similar sensitivity node for LLO. About 15% of injected signals have a 50% difference or larger in the effective distance between the sites due to the slight misalignment of the detectors.

the H1 trigger is greater than the maximum distance to which H2 can detect a trigger we keep the H1 trigger without searching for coincidence.

If coincident triggers are found in H1 and H2, we can get an improved estimate of the amplitude of the signal arriving at the Hanford site by coherently combining the filter outputs from the two gravitational-wave channels,

$$\rho_H = \sqrt{\frac{|z_{H1} + z_{H2}|^2}{\sigma_{H1}^2 + \sigma_{H2}^2}}. \quad (9)$$

The more sensitive interferometer receives more weight in this combination, as can be seen from Eq. (3), in which the noise enters in the denominator. If a trigger is found in only one of the Hanford interferometers, then ρ_H is simply taken to be the value of ρ from that interferometer.

The final step of the search is to apply the POB_I veto, described in Sec. V, to eliminate certain L1 triggers which arose from instrumental glitches. The individual signal-to-noise ratios are used to construct a multidetector statistic as described in Sec. VIII for any surviving triggers.

The surviving coincident triggers are clustered in a way that identifies the best parameters to associate with a possible inspiral signal in the data. The clustering is needed since large astrophysical signals and instrumental noise bursts can produce many triggers with coalescence times within a few seconds of each other. We chose the trigger with the largest SNR from each cluster; triggers separated

by more than 4 seconds were considered unique. Alternative clustering methods are discussed in Sec. VII in an effort to understand the accidental likelihood of a small number of candidates that were observed in the final sample.

To perform the search on the full data set, a directed acyclic graph (DAG) was constructed to describe the work flow, and execution of the pipeline tasks was managed by Condor [31] on the UWM and LIGO Beowulf clusters. The software to perform all steps of the analysis and construct the DAG is available in the package LALAPPS [30].

B. Parameter tuning

The entire analysis pipeline was studied first using the playground data set to tune the values of the various parameters. The goal of tuning was to maximize the efficiency of the pipeline to detection of gravitational waves from binary inspirals without producing an excessive rate of spurious candidate events. The detection efficiency was determined by Monte-Carlo simulations in which we injected simulated inspiral signals from our model population into the data. The efficiency is the ratio of the number of signals detected to the number injected, as described in Sec. VI B 1. In the absence of a detection, a pipeline with a high efficiency and low false alarm rate allows us to set a better upper limit, but it should be noted that our primary motivation is to enable reliable detection of gravitational waves. By evaluating the efficiency using Monte-Carlo injection of signals from the hypothetical population of binary neutron stars into the data, we account for systematic effects caused by our vetoes and other aspect of our pipeline.

There are two sets of parameters that we can tune in the pipeline: (1) the single interferometer parameters which are used in the matched filter and χ^2 veto to generate inspiral triggers in each interferometer, and (2) the coincidence parameters used to determine if triggers from two interferometers are coincident. The single interferometer parameters include the signal-to-noise threshold ρ^* , the number of frequency sub-bands p in the χ^2 statistic, the coefficient δ on the SNR dependence of the χ^2 cut, and the χ^2 cut threshold ξ^* . These are tuned on a per-interferometer basis, although some of the values chosen are common to two or three detectors. The coincidence parameters are the time-coincidence window for triggers, δt , the mass parameter coincidence window δm , and the effective distance cut parameters ϵ and κ described in Eq. (8). Because of the nature of the triggered search pipeline, parameter tuning was carried out in two stages. We first tuned the single interferometer parameters for the primary detector (L1). We then used the triggered template banks (generated from the L1 triggers) to explore the single interferometer parameters for the less sensitive Hanford detectors. Finally, the parameters of the coincidence test were tuned.

1. Single interferometer tuning

The number of bins p used in the χ^2 test was set to $p = 15$ (compared with $p = 8$ used in our S1 search [5]), in order to better differentiate spurious triggers from actual (or injected) signals, while still having at least several cycles of the waveform in each bin.

The δ parameter in Eq. (7) is expected to be no less than 0.03 for our choice of maximum 3% SNR loss in the L1 template bank. A dedicated investigation using software injections into L1 showed that the χ^2 test rejected 50% of signals from the Milky Way (with effective distances closer than 200 kpc) when using $\delta = 0.03$; using $\delta = 0.1$ recovered all such injections. The search efficiency for weaker injections, at effective distances larger than 900 kpc, did not depend on δ .

The signal-to-noise threshold $\rho^* = 6$ was used in all three instruments. This choice was motivated by the observation that a signal, with certain orbital orientations and sky positions, can have a smaller effective distance in the less sensitive detector (see Sec. VIB 2 and Fig. 6). This choice of threshold was computationally possible due to

TABLE II. The effect of lowering the χ^2 threshold, ξ^* , for the single interferometer L1, and for the combination of L1 and H1 in the pipeline. The efficiency of the L1 search remains constant as the threshold is lowered to 5.0; however, as the threshold is lowered in H1, the efficiency of the triggered search drops. Further testing indicated that a threshold of 12.5 in H1 was acceptable without a loss of efficiency.

Value of χ^2 threshold, ξ_{L1}^*	L1 efficiency	Pipeline efficiency
20.0	0.350	0.270
15.0	0.350	0.270
10.0	0.350	0.255
5.0	0.350	0.212

TABLE III. A complete list of the parameters that were selected at the various stages of the pipeline. The procedures used to select these parameter values are outlined in the text.

Parameter	Pipeline node	Value
ρ^*	GW data (all)	6.0
p	GW data (all)	15
δ	GW data (all)	0.1
ξ_L^*	L1 GW data	5.0
ξ_H^*	H1/H2 GW data	12.5
δt_{HH}	H1/H2 trigger coincidence	0.001 s
δt_{HL}	L1/H1 and L1/H2 trigger coincidence	0.011 s
δm	Trigger coincidence (all)	0.0
κ_{HH}	H1/H2 trigger coincidence	0.5
ϵ_{HH}	H1/H2 trigger coincidence	2
κ_{HL}	L1/H1 and L1/H2 trigger coincidence	1000

our efficient pipeline and code; it also provided better statistics for background estimation (Sec. VII).

Once the L1 search parameters had been tuned, the resulting triggered template banks were used as input to tune the H1 and H2 χ^2 threshold parameters ξ^* in the coincidence search. The final value was selected so that the triggered search suffered no loss of efficiency due to this single parameter. Table II shows how the parameter ξ^* was tuned, first for L1 and then for H1. The values chosen for all parameters are shown in Table III.

2. Coincidence parameter tuning

After the single interferometer parameters had been selected, the coincidence parameters were tuned. As described in Sec. VIC, the coalescence time of an inspiral signal can be measured to within ≤ 1 ms. The gravitational waves' travel time between observatories is 10 ms, so δt was chosen to be 1 ms for LHO-LHO coincidence and 11 ms for LHO-LLO coincidence. The mass coincidence parameter was initially chosen to be $\delta m = 0.03$; however, testing showed that this could be set to $\delta m = 0.0$ (i.e. requiring the triggers in each interferometer to be found with the exact same template) without loss of efficiency.

After tuning the time and mass parameters, we tuned the effective distance parameters κ and ϵ . Initial estimates of $\epsilon = 2$ and $\kappa = 0.2$ were used for testing; however, we noticed that many injections were missed when testing for LLO-LHO distance consistency. This is due to the slight detector misalignment between the two sites from Earth's curvature, which causes the ratio of effective distances, as measured at the two observatories, to be large for a significant fraction of our target population, as shown in Fig. 6. Consequently, we disabled the effective distance consistency requirement for triggers generated at different observatories. A study of simulated events injected into H1 and H2 suggested values of $\epsilon_{HH} = 2$ and $\kappa_{HH} = 0.5$ to be suitable. Note that, as described above, we demand that an L1/H1 trigger pass the H1/H2 coincidence test if the effective distance of the trigger in H1 is within the maximum range of the H2 detector at threshold.

C. Validation

Hardware signal injections were used as a test of our data analysis pipeline. These injections allow us to study issues of instrumental timing and calibration, as well as to verify that the injected signals were indeed identified as triggers in our analysis pipeline. At intervals throughout S2, a predetermined set of inspiral and burst signals were injected into the instruments using the mirror actuators.

We examined six sets of hardware injections spread throughout the S2 science run. Each set had six 1.4–1.4 M_{\odot} inspiral signals at effective distances spaced logarithmically between 500 kpc and 15 kpc, and four 1.0–1.0 M_{\odot} inspirals at distances from 500 kpc to 62 kpc.

Each strain waveform was calculated using the second order post-Newtonian expansion for an optimally oriented inspiral, and appropriately scaled for an inspiral at the desired effective distance. The strain was then converted into an injection signal into the y arm of the interferometer, with the appropriate calibration to produce the desired differential strain.

Of these injections, five sets were injected simultaneously into all three instruments and the sixth was injected into L1 and H1 only. After the coincidence stage of the pipeline, 59 of the 60 injections produced a trigger in the appropriate template ($1.4\text{--}1.4M_{\odot}$ or $1.0\text{--}1.0M_{\odot}$). (The single missed injection had a χ^2 value in L1 which was slightly above threshold.) This is an excellent test of every stage of our pipeline. L1 successfully produced triggers corresponding to the injections, which were then used to make triggered banks against which the Hanford instruments were analyzed. Both H1 and H2 produced triggers which were coincident with those in L1. Furthermore, the amplitude cut between H1 and H2 was effective. For the louder injections, the effective distance was consistent between the two instruments. In the case of more distant injections which did not produce triggers in H2, the coincidence stage of the pipeline allowed the L1-H1 coincident triggers to be kept even though no event was found in H2.

The pipeline was also run on an additional set of hardware injections at distances between 5 Mpc and 150 kpc. Of these injections, the most distant which produced coincident triggers were at 1.25 Mpc for the $1.4\text{--}1.4M_{\odot}$ injection and 620 kpc for the $1.0\text{--}1.0M_{\odot}$ injection. This is consistent with the sensitivity of H1 during the injection time, and with the efficiency measured in our software injections (Table I).

The trigger times associated with the hardware injections all agreed with expectations to within 1 msec.

Hardware injections provide a powerful test of instrumental calibration. Any errors in the calibration of the instrument affect the measured effective distance to the inspiral. Uncertainties in distance measurements will contribute to errors in our detection efficiency. Additionally, they become important when requiring consistency of effective distance in triggers from the two Hanford instruments. To test the accuracy of the effective distance measurements, we used the loudest injections—namely the $1.4\text{--}1.4M_{\odot}$ injections with distances less than 100 kpc—so that systematic errors in distance measurement would dominate over noise. The effective distance of all such injections was accurate to within 20%. We found that for L1, the distance to injections was systematically underestimated by 5%, with a 4% standard deviation. For H1, there was a systematic underestimation of 2% with a standard deviation of 3%. For H2, we had a 2% overestimation and a standard deviation of 5%. These errors are consistent with uncertainties in the calibration described in Sec. . Further details of the S2 hardware injections are available in [32].

A Markov chain Monte-Carlo routine (MCMC) [33] was also used to examine the injected signals. This provides a method of estimating the parameters of an injected signal. As an example, for a $1.4\text{--}1.4M_{\odot}$ injection at 125 kpc, the MCMC routine generated masses of $1.4003M_{\odot}$ and $1.3991M_{\odot}$. The width of the 95% confidence interval for each mass was $0.012M_{\odot}$.

VII. BACKGROUND ESTIMATION

An event candidate which survives all cuts in our analysis pipeline, including coincidence, can arise either from a real gravitational-wave signal or from noise bursts which contaminate our data streams. We refer to the latter class of event candidates as background. These background events are caused by many different environmental and instrumental processes. Under the assumption that such processes are uncorrelated between the detectors at Hanford and Livingston, we estimated the rate for background events due to accidental coincidences by applying artificial time shifts Δt to the triggers coming from the Livingston detector. These time-shift triggers were then fed into subsequent steps of the pipeline. For a given time shift, the triggers that survived to the end of the pipeline represent a single trial output from our search (if no coincident gravitational-wave signals were present).

A total of 40 time shifts were analyzed to estimate the background: $\Delta t = \pm 5, \pm 10, \pm 15, \pm 27, \pm 37, \pm 47, \pm 57, \pm 67, \pm 77, \pm 87, \pm 97, \pm 107, \pm 117, \pm 127, \pm 137, \pm 147, \pm 157, \pm 167, \pm 177, \pm 197$ seconds. To avoid correlations, we used time shifts longer than the duration of the longest template waveform (~ 4 seconds). We did not time shift the triggers from the Hanford detectors relative to one another since real correlations may arise from environmental disturbances. The resulting distribution of time-shift triggers in the (ρ_H, ρ_L) plane was used to determine a joint signal-to-noise ρ . Here ρ_L is given by Eq. (5) and ρ_H is given by Eq. (9). For Gaussian noise fluctuations, with the single interferometer triggers' SNR maximized over the polarization phase and the orbital inclination angle, one expects circular false alarm contours centered on the origin [34], suggesting that the sum of the squares of the signal-to-noise ratios would be a useful combined statistic. The time shifts revealed the need for a modified statistic; we settled on contours of (roughly) constant false-alarm probability to use in assigning a signal-to-noise to coincident triggers. As shown in Fig. 7, the combined SNR

$$\rho = \sqrt{\rho_L^2 + \rho_H^2/4} \quad (10)$$

yielded approximate constant-density contours for the distribution of background events in the plane.

Using data from the time shifts, we also computed the (sample) mean number of events per S2 with $\rho > \rho^*$. This




FIG. 7 (color online). The signal-to-noise at Livingston ρ_L plotted against the signal-to-noise at Hanford ρ_H for triggers from 40 time shifts. In Gaussian noise, one expects circular false alarm contours centered on the origin indicating that the sum of the squares of the signal-to-noise ratios would be a useful combined statistic. The observed distribution of time-shifted triggers motivated the modified statistic presented in Eq. (10).

result is shown in Fig. 8; the shaded bars represent the sample variance for ease of comparison with the zero-shift distribution. The apparent exponential dependence of $b(\rho^*)$ on ρ^{*2} further supports the choice of combined statistic in Eq. (10).




FIG. 8. The number of triggers per S2 above combined SNR ρ^* . The triangles represent the expected (mean) background based on 40 time-shift analyses. The shaded envelope indicates the sample variance in the number of events. The choice of combined SNR in Eq. (10) is further justified by the behavior $b(\rho^*) \propto \exp(-\rho^{*2})$ down to about 0.3 events per S2. The circles represent the unshifted inspiral event candidates (Sec. VIII A).

VIII. SEARCH RESULTS

The pipeline described above was used to analyze the S2 data. The output of the pipeline is a list of candidate coincident triggers. To decide whether there is any plausible detection candidate worth following up, we compare the combined SNR of the candidates with the expected SNR from the accidental background. If the probability of any candidate being accidental is small enough, we look at the robustness of the parameters of the candidate under changes in the pipeline, and we investigate possible instrumental reasons for these candidates that may have been overlooked in the initial analysis.

Independently of whether detection candidates are found, we can use the results to set an upper limit on the rate of binary neutron star coalescences per Milky Way equivalent galaxy (MWEG), per year. We use the same statistics as in the previous search in S1 data [5], measuring the efficiency of the search at the SNR of the loudest trigger found. We take here the most conservative approach, taking into account *all* triggers at the output of the pipeline: both those triggers considered to be potential detection candidates *and* those that are consistent with being due to instrumental noise or consistent with background. We do not include the playground in the observation time used for calculating the upper limit, since the playground was used to tune the pipeline. This is consistent with our approach of focusing on detection, and not optimizing the pipeline for upper limit results (which was the reason for not considering single detector data, for example).

As described in Sec. II, after the data quality cuts, discarding science segments with durations shorter than 2048 s, and application of the instrumental veto in L1, a total of 373 hours of data were searched for signals, broken up in double- and triple-coincidence time as shown in Fig. 2. For the upper limit analysis, including the background estimation, we only considered the nonplayground times, amounting to 339 hours, of which 65% (221 hours) had all three detectors in operation; 26% (89 hours) had only L1 and H1 in operation; and 8.5% (29 hours) had only L1 and H2 in operation.

A. Triggers and event candidates

The output of the pipeline is a list of candidates which are assigned a SNR according to Eq. (10). There are 142 candidates in the nonplayground final sample with combined ρ^2 greater than 45; the breakdown is 90 candidates from L1-H2 two-detector data, 35 from triple-detector data, and 17 from L1-H1 two-detector data. All the candidates in the triple-detector data had SNR in H1 too small to cross the threshold in H2, and following our pipeline, were accepted as coincident triggers. Thus, all our coincident triggers are only double-coincidence candidates.

Table IV lists the ten largest SNR coincident triggers recorded in the analysis (including the playground).

TABLE IV. The 10 triggers with the largest SNR which remain at the end of the pipeline. This table indicates their UTC date, the time delay between Hanford and Livingston ($\delta t = t_H - t_L$), the combined SNR ρ^2 [from Eq. (10)], the SNR registered in each detector, the value of χ^2 per degree of freedom at each interferometer (for $2p - 2 = 28$ d.o.f.), the effective distance to an astrophysical event with the same parameters in each detector, and the binary component masses of the best matching template (identical for the triggers in both detectors). The notation L1-Hx means that all three interferometers were in science mode at the time of this coincident trigger.

Rank	YYMMDD (UTC)	δt (ms)	Instruments	ρ^2	ρ_L	ρ_H	$\chi^2/\text{d.o.f. (LLO)}$	$\chi^2/\text{d.o.f. (LHO)}$	D_{eff}^L (Mpc)	D_{eff}^H (Mpc)	$m_1 (M_\odot)$	$m_2 (M_\odot)$
#1	030328	+9.8	L1-H2	89.1	8.9	6.1	2.6	3.4	2.1	0.9	2.23	2.23
#2	030327	-1.7	L1-H2	87.1	7.6	11.0	2.7	5.1	2.1	0.7	2.09	2.09
#3	030224	-7.6	L1-H2	79.5	6.5	12.3	1.2	6.4	2.6	0.4	3.04	1.19
#4	030301	+7.1	L1-H1	79.0	7.7	8.8	2.3	4.7	2.1	0.6	1.45	1.45
#5	030224	+5.4	L1-H1	77.2	6.7	11.3	1.9	6.5	2.0	0.5	1.90	1.35
#6	030301	+11.0	L1-H1	69.2	6.7	9.9	2.0	5.9	2.2	0.5	1.28	1.28
#7	030224	-1.5	L1-H2	67.2	6.9	8.9	1.3	4.1	2.7	0.8	2.46	1.82
#8	030328	+4.4	L1-H2	67.1	7.4	7.2	2.8	2.3	1.9	0.6	2.27	1.11
#9	030224	-10.5	L1-H1	66.5	6.9	8.8	1.6	4.3	2.3	0.8	3.60	1.20
#10	030301	-1.7	L1-Hx	65.9	7.5	6.1	2.5	2.7	2.5	1.6	2.65	2.65

Detailed investigations of the conspicuous triggers were performed and are reported below. Nine out of these ten candidates were in the double-detector sample. There was only one (#10 in the list) which was in the triple-detector data, although it was too weak to require a corresponding trigger in H2. Our loudest candidate, as well as five of our loudest ten, are L1-H2 coincident triggers; in fact, 63% of all our candidates are L1-H2 coincident triggers. Since the L1-H2 data supplies only 9% of the total observation time, we conclude that the noise in H2 was significantly different than the noise in H1, producing more and louder triggers. For all triggers in Table IV, the effective distance is larger in the Livingston detector than in the Hanford detector: although this is plausible for real signals, as shown in Fig. 6, it suggests that these candidates are more likely to originate from instrumental noise.

A trigger gets elevated to the status of an event candidate if the chance occurrence due to noise is small as determined by time-shift background estimation, calculated as described in Sec. VII. Each event candidate is subjected to follow-up investigations beyond the level of automation used in our pipeline to ensure that it is not due to an instrumental or environmental disturbance.

In Fig. 8, a cumulative histogram of the final coincident triggers versus ρ^2 is overlaid on the expected background due to accidental coincidences, as determined by time shifts. Even after taking into account that these are cumulative histograms, so that adjacent bins are strongly correlated, it appears that the number of coincident triggers is inconsistent with the expected background for large ρ^2 . While the origin of this discrepancy is not understood, careful examination of the three coincident triggers with the highest ρ^2 , detailed below, indicates that these are not gravitational-wave detections. Moreover, no evidence of correlated noise between the Livingston and Hanford observatories was found for times around these triggers. In Sec. VIII B, we demonstrate that a reasonable modification

of our algorithm for clustering multiple coincident triggers results in good agreement between the coincident trigger sample and the background estimate. This shows that the background estimate is not robust with respect to reasonable variations in the analysis procedure.

In Sec. IX we derive an upper limit on the rate of inspiral events which is conservative with respect to any uncertainties in the background estimate. Thus, although the discrepancy between the coincident triggers and the background estimate presented in Fig. 8 is not understood, it does not affect our ability to detect inspiral signals or set an upper limit on the inspiral rate.

1. Trigger 030328

The loudest candidate is in a cluster of three coincident triggers in L1 and H2 on March 28, 2003. The loudest trigger in the cluster, #1 in Table IV, has a large SNR in L1 ($\rho = 8.9$), but a SNR in H2 ($\rho = 6.1$) which is close to threshold for trigger generation. We expect a candidate arising from a real signal to be robust under small changes in our analysis pipeline. In order to test the robustness of this particular candidate upon changes in the boundaries for the blocks used in the analysis, we reanalyzed the data in this science segment shifting the start time of the block by different amounts in L1 and H2. Although in all cases our analysis produced similar numbers of triggers in each detector, only in the original case were there triggers within 11 ms that had identical masses and were considered candidates.

Other measured parameters of this trigger reinforce the conclusion that it should not be promoted to a detection candidate. The χ^2 per degree of freedom is bad, especially in H2 (3.4); the χ^2 in L1 is close to our threshold [$\chi^2/(p + \delta^2 \rho^2) = 4.6$ compared to the threshold $\xi^* = 5$]. The effective distances measured by the detectors are very different, 2.1 Mpc in L1 and 0.9 Mpc in H2. This ratio of

effective distances and the large time delay (10 ms) is unlikely, though not impossible, for the posited population. The other triggers in the cluster (with smaller SNR by definition) have a larger distance ratio, smaller time delay, smaller χ^2 , and different masses in the binary system (but the same masses in both triggers): $2.9M_{\odot}$ and $1.1M_{\odot}$.

The time series in the H2 detector does not show anything unusual in the gravitational-wave channel or in other auxiliary channels: this is consistent with the small SNR measured in H2 for this trigger. The time series in L1, however, shows several disturbances in a few second window around the candidate. Many inspiral triggers are generated at the time, although only three are coincident in mass and time with H2 triggers. Just 100 ms before the coalescence time of the trigger in question (and within the duration of the template corresponding to the trigger), we observed a fast fluctuation in the calibration, less than 62 ms in duration, indicating low optical gain dropping to levels that would make the feedback loop unstable. In the 51 minutes when the L1 detector was operational around the time of the candidate, such a low loop gain happened at three different instances. Each occurrence lasted less than 62 ms and was accompanied by a cluster of inspiral triggers generated by the search code. Since we averaged the calibration on a 60-second time scale, these rare fluctuations had not been observed when considering data quality. Nevertheless, such a low gain, even though brief, makes the data and its calibration very unreliable. This provides an instrumental reason to veto this trigger as a detection candidate: the cluster of triggers in L1 containing the coincident trigger was probably produced by non-linear effects associated with either the low gain itself or the alignment fluctuation that produced the low gain in the first place. An indication of this is that most of the excess power at the time of the trigger cluster is at sidebands of a narrow line in the spectrum, corresponding to up-conversion of low-frequency noise to the 120 Hz electric power line harmonic by some unknown bilinear process.

2. Trigger 030327

The second loudest candidate in Table IV is on March 27, 2003; it is also an L1-H2 trigger, with a combined SNR $\rho^2 = 87$. Similar to the loudest candidate, according to the background estimates, it has a 20% false alarm probability. Unlike the loudest candidate, however, this trigger has a smaller SNR in L1 (7.6) than in H2 (11.0).

As shown in Fig. 9, there are four bursts of noise in the gravitational-wave channel of the H2 detector in the 58 minutes surrounding this candidate when L1 and H2 were operating in coincidence. Each burst is a few to ten seconds long, and each triggers a large fraction of the template bank. Trigger 030327 is the loudest trigger in a cluster of seven L1-H2 coincident triggers, in a 3 ms interval, during one of the H2 noise bursts. There are also three other coincident triggers in the same noise burst, and there are




FIG. 9 (color online). Triggers produced by the inspiral search in L1 (top panel), and the triggered search in H2 (middle panel) and H1 (lower panel), near the time of our second loudest candidate, on 03/03/27 UTC, in Table IV. The black points indicate the triggers that have identical mass and are coincident within a 11 ms window. The circles indicate the clustered candidate triggers. The dashed lines are the boundaries of the science segments in each detector (i.e., H1 and H2 were not operating to the left of the dashed line in their respective graphs, and L1 was not operating after the dashed line in the top panel).

clusters of coincident triggers in two of the other three H2 noise bursts. One of these bursts, with several coincident triggers, contains another one of our top 20 loudest candidates. The only noise burst that does not have any coincident triggers happens at a time when the H1 detector had begun operating, and the triple-coincidence criteria were not satisfied. In almost all of the time shifts used to estimate the background, there were a number of L1-H2 coincident triggers at the time of the noise bursts in H2. We have not found the precise instrumental origin of the disturbances causing such large numbers of inspiral triggers: although simultaneous disturbances are observed in a few other auxiliary channels, none of them are as obvious as they are in the gravitational-wave channel.

Despite a relatively low measured probability of this trigger being due to the background, we do not consider this trigger an event candidate due to the low SNR in the L1 detector, the poor χ^2 , the unlikely parameters, and a strong suspicion of instrumental misbehavior.

3. Trigger 030224

We also followed up the third loudest candidate which occurred on February 24, 2003. This is again an L1-H2 candidate, happening a few minutes before H2 turns off due to loss of the operating point at optical resonance. We discovered a very large feedback loop oscillation in an auxiliary servo operating in the orthogonal channel to the gravitational channel (L1:LSC-AS_I). Although the oscil-

lation was at lower frequencies than the ones relevant for our search, it produced broadband excess power at all frequencies in two short time intervals, at the beginning and at the end of the oscillation. This instrumental misbehavior rules out this trigger as an event candidate.

B. Background estimation revisited

One striking feature of our analysis is the apparent discrepancy between the coincident trigger count in the zero-shift data set compared to our background estimate shown in Fig. 8. Our follow-up investigations rule out gravitational-wave signals as the origin of this discrepancy. Moreover, no evidence of correlated noise between the Livingston and Hanford observatories was found for times around these triggers. We thus expect that the loudest triggers should be consistent with a background caused by instrumental artifacts, but Fig. 8 suggests otherwise. In an effort to understand how selection effects in our pipeline might be responsible for this apparent discrepancy, we repeated our analysis with a different method of clustering the coincident triggers at the end of our pipeline. Instead of selecting the trigger with the largest SNR from each cluster, we selected the coincident trigger with the minimum value of

$$\frac{\chi_H^2}{p + \delta^2 \rho_H^2} + \frac{\chi_L^2}{p + \delta^2 \rho_L^2}. \quad (11)$$

[Compare with Eq. (7), and cf Eq. (10).] We call this *best-fit clustering* since it preferentially selects events with small χ^2 and/or large SNR. (See Sec. VIA for a description of the maximum SNR clustering.) The resulting background estimate and zero-shift distribution are shown in Fig. 10; no evidence of the original discrepancy remains.

In order to estimate the effect of this new clustering method in the efficiency of our search, we show in Fig. 11 a comparison of the SNR under each clustering scheme for the final S2 trigger sample and a simulated injection run. The plot shows that, as expected, the SNR from best-fit clustering is either the same or smaller than the corresponding SNR using maximum SNR clustering. Moreover the values of ρ^2 associated with triggers 030328 and 030327, the two loudest triggers in Table IV, are less than 55 for best-fit clustering. Simulated injections, however, produce similar values of ρ^2 for both clustering methods. We see then that the best-fit clustering does not hurt detection efficiency (injections retain their SNR), and makes identified instrumental artifacts consistent with the estimated background (specifically, triggers 030328 and 030327). More importantly, in our opinion, it indicates that real signals would be robust under changes of clustering, thus adding further weight to our conclusion that the two loudest triggers in Table IV are *not* gravitational waves from inspiral signals.

Although the clustering method described in this section is deemed better for detection purposes, we will use our




FIG. 10. The number of triggers per S2 above combined SNR ρ^* using the best-fit clustering method. The triangles represent the expected (mean) background, while the circles represent zero-shift coincident triggers. See Fig. 8 and Sec. VII for details of the time shifts and for comparison with largest SNR clustering. We note that there is no apparent excess of S2 coincident triggers over the expected background from accidental coincidences in this plot.

original clustering method as described in Sec. VIA for an upper limit on the rate of coalescences per galaxy, to avoid possible biases.




FIG. 11. A comparison between the SNR associated with a cluster of triggers by maximizing the SNR over the cluster and the SNR associated with the same cluster using best-fit clustering. Notice that the best-fit clustering often gives lower SNR values for the final S2 sample, whereas simulated injections have very similar SNR values for both clustering methods. The triggers 030327, 030328, and 030224 discussed in the text are the gray circles with horizontal coordinates 89.1, 87.1, and 79.5, respectively.

IX. UPPER LIMIT ON THE RATE OF COALESCENCES PER GALAXY

Following the notation used in [5], let \mathcal{R} indicate the rate of binary neutron star coalescences per year per MWEG and $N_G(\rho^*)$ indicate the number of MWEGs to which our search is sensitive at $\rho \geq \rho^*$. The probability of observing an inspiral signal with $\rho > \rho^*$ in an observation time T is

$$P(\rho > \rho^*; \mathcal{R}) = 1 - e^{-\mathcal{R}TN_G(\rho^*)}. \quad (12)$$

A trigger can arise from either an inspiral signal in the data or from background. If P_b denotes the probability that all background triggers have SNR less than ρ^* , then the probability of observing one or more triggers with $\rho > \rho^*$ is given by

$$P(\rho > \rho^*; \mathcal{R}, b) = 1 - P_b e^{-\mathcal{R}TN_G(\rho^*)}. \quad (13)$$

Given the probability P_b , the total observation time T , the largest observed signal-to-noise ρ_{\max} , and the number of MWEGs $N_G(\rho_{\max})$ to which the search is sensitive, we find that the rate of binary neutron star inspirals per MWEG satisfies

$$\mathcal{R} < \mathcal{R}_{90\%} = \frac{2.303 + \ln P_b}{TN_G(\rho_{\max})} \quad (14)$$

with 90% confidence. This is a frequentist upper limit on the rate. For $\mathcal{R} > \mathcal{R}_{90\%}$, there is more than 90% probability that at least one event would be observed with SNR greater than ρ_{\max} . Details of this method of determining an upper limit can be found in Ref. [35]. In particular, one obtains a conservative upper limit by setting $P_b = 1$; we adopt this approach below because of uncertainties in our background estimate.

During the $T = 339$ h of data used in our analysis, the largest observed SNR was $\rho_{\max}^2 = 89.1$. The number of MWEGs N_G was computed using a Monte-Carlo simulation in which the data was reanalyzed with simulated inspiral signals, drawn from the population described in Sec. III, added to the time series. The results are shown in Fig. 12 which breaks down the contribution to N_G galaxy by galaxy in the target population. (Results are shown only for galaxies contributing more than 0.01 MWEG.) At $\rho^{*2} = 89.1$, we find $N_G = 1.34$ MWEG; this is subject to some uncertainties, to be discussed in the next section. We determine that the probability that all background events have SNR smaller than the largest observed SNR is $P_b = 0.8$. Since no systematic error has been assigned to this background estimate, we take $P_b = 1$ to be conservative. As a function of the true value of N_G , the rate limit is

$$\mathcal{R}_{90\%} = 44.4 \left(\frac{1.34}{N_G} \right) \text{yr}^{-1} \text{MWEG}^{-1}. \quad (15)$$




FIG. 12 (color online). The efficiency of the search to the target population described in Sec. III as a function of the SNR threshold ρ^* . The contribution of each galaxy to N_G is shown for galaxies which contribute more than 1% of a MWEG. The largest SNR of a coincident trigger observed in this analysis was $\rho^2 = 89$ meaning that $N_G = 1.34$ MWEG were probed by the search.

A. Error analysis

The principal systematic effects on our rate limit are (i) inaccuracies in our model population, including inaccuracies in the inspiral waveform assumed, and (ii) errors in the calibration of the instrument. All other systematic effects in the analysis pipeline (for example, less-than-perfect coverage of the template bank) are taken into account by the Monte-Carlo estimation of the detection efficiency. It is convenient to express the effective number of MWEGs as

$$N_G = \varepsilon(\rho_{\max}) \left(\frac{L_{\text{pop}}}{L_{\text{MW}}} \right) \quad (16)$$

$$= \sum_{i=0}^n \varepsilon_i(\rho_{\max}) \frac{L_i}{L_{\text{MW}}} \quad (17)$$

where L_{MW} is the effective blue-light luminosity of the Milky Way and the index i identifies a galaxy in the population where $i = 0$ corresponds to the Milky Way. The fraction of the signals that would be detectable from a particular galaxy in this search is denoted ε_i —we will refer to this as the efficiency of the search to sources in galaxy i . Referring to Table I, $\varepsilon_i(\rho_{\max})$ is given in column 6, L_i/L_{MW} in column 5, and the cumulative value of N_G in column 7.

1. Uncertainties in population model

Uncertainties in the population model used for the Monte-Carlo simulations may lead to differences between the inferred rate and the rate in the Universe. Since the

effective blue-light luminosity L_{pop} is normalized to our Galaxy, variations arise from the relative contributions of other galaxies in the population. These contributions depend on the estimated distances to the galaxies, estimated reddening, and corrections for metallicity (lower values tend to produce higher-mass binaries), among other things.

The spatial distribution of the sources can also introduce significant uncertainties. Typically, the distances to nearby galaxies are only known to about 10% accuracy. In fact, the distance to Andromeda is thought to suffer from a 15% error. A change in distance Δd to a galaxy assumed to be at distance d introduces two different uncertainties in N_G :

- (1) The change in our estimate of the efficiency for the particular galaxy is determined by observing that the correct signal-to-noise ratio decreases inversely with the distance to the galaxy and hence

$$\varepsilon_i(\rho_{\text{max}}) \rightarrow \varepsilon_i(\rho_{\text{max}}[1 + \Delta d/d]) \quad (18)$$

so that efficiency decreases as the distance increases.

- (2) The change in our estimate of the luminosity that is correlated with this change in distance is given by

$$\frac{L_i}{L_{\text{MW}}} \rightarrow \frac{L_i}{L_{\text{MW}}} \left(\frac{d + \Delta d}{d} \right)^2 \quad (19)$$

so that our estimate of the luminosity increases as the distance increases.

Adopting a 15% error in all distances, we estimate the error in N_G due to uncertainties in galactic distances to be

$$\delta N_G|_d = 0.04 \text{ MWEg}. \quad (20)$$

The absolute blue-light luminosity of the Milky Way is uncertain since it is inferred from other galactic parameters. There appears to be some confusion about this in the literature. Phinney [11] uses $L_{\text{MW}} = 1.6 \times 10^{10} L_{\odot,B}$ while the more recent work of Kalogera, Narayan, Spergel, and Taylor [36] uses $9 \times 10^9 L_{\odot,V}$ citing a factor of 2 difference with the number used by Phinney; these numbers agree within 1%, however, once converted to the same units.

Errors in the blue magnitude of each galaxy and relative normalization to the Milky Way enter in the same way. For the galaxies that contribute most substantially to N_G , the Lyon-Meudon Extragalactic Database catalog suggests errors of $\sim 10\%$ in corrected blue magnitude which translates into errors of $\sim 20\%$ in luminosity. Adopting this uniformly for the galaxies in our population, we estimate

$$\delta N_G|_L = 0.04 \text{ MWEg}. \quad (21)$$

Different models for BNS formation can lead to variations in the neutron star mass distribution [12], although the bulk of the distribution always remains strongly peaked around observed neutron star masses [37]. We estimate the corresponding variations in N_G to be

$$\delta N_G|_{\text{mass}} \approx 0.01 \text{ MWEg} \quad (22)$$

based on simulations with 50% reduction in the number of binary systems with masses in the range $1.5M_{\odot} < m_1, m_2 < 3.0M_{\odot}$. The error estimate is only indicative; a few extreme scenarios for binary formation can produce more severe alterations to the mass distribution [12].

The waveforms used both in the Monte-Carlo simulation and in the detection templates ignore spin effects. Estimates based on the work of Apostolatos [18] suggest that less than 10% of all spin orientations and parameters consistent with binary neutron stars provide a loss of SNR greater than $\sim 5\%$. Thus, the mismatch between the signal from spinning neutron stars and our templates should not significantly affect the upper limit. To be conservative, however, we place an upper limit only on nonspinning neutron stars; we will address this issue quantitatively in future analysis.

2. Uncertainties in the instrumental response

The uncertainty in the calibration can be quantified as an uncertainty in amplitude and phase of a frequency-dependent response function.

Using the same procedure as in S1 [38], the average instrument response $R(f)$ was constructed for every minute of data during S2 from a reference sensing function $C(f)$, a reference open loop gain function $G(f)$, and a parameter $\alpha(t)$ representing varying optical gain. The response function at a time t during the run is given by $R(f, t) = (1 + \alpha(t)G(f))/(\alpha(t)C(f))$. The parameter α was reconstructed using the observed amplitude of a calibration line. If an inspiral signal is present in the data, errors in the calibration can cause a mismatch between the template and the signal. We can quantify and measure this effect using simulated injections, both in software and hardware. For such injections, the SNR differs from the SNR that would be recorded for a signal from a real inspiral event at the same distance as the injection. The effect is linear in amplitude errors causing either an upward or downward shift in SNR; in contrast, the effect is quadratic in phase errors causing an overestimation of sensitivity [39]. This error is propagated to our upper limit by shifting the efficiency curve in Fig. 12 horizontally by the appropriate amount.

A careful evaluation of uncertainties in the S2 calibration [40] has shown that amplitude errors have two dominant components. The first source of error, an imperfect knowledge of the strength of the feedback actuators in the system, produces an error in the overall amplitude of the sensing function $C(f)$. Thus, it manifests as a systematic error in the amplitude of the response function, constant during the whole run. This error was estimated to be 8.5% in L1, 3.5% in H1, and 4.5% in H2.

The second dominant source of error is an imprecise measurement of the amplitude of the calibration line, resulting in an error of the coefficient $\alpha(t)$. This measurement error is mostly random in nature, and leads to mag-

nitude and phase errors in the response function. These errors translate into maximum amplitude errors in the response function equal to 6% for L1 and H2, and 18% for H1.

Another source of error in $\alpha(t)$, which is not captured in the errors as estimated above, is due to changes in optical gain that are averaged over the 60 s integration time used in the measurement of $\alpha(t)$. Based on a limited set of diagnostics, we roughly estimate these errors can be as large as 10% in L1, and smaller in H1 and H2. The errors in α , including these fast fluctuations, are random in nature, and are not expected to contribute to the error in the measured efficiency.

The SNR used in the final analysis is constructed from the individual SNR's ρ_L and ρ_H as given in Eq. (10). The error in SNR then has two pieces,

$$(\delta\rho)^2 = \frac{\rho_L^2}{\rho^2}(\delta\rho_L)^2 + \frac{1}{4^2} \frac{\rho_H^2}{\rho^2}(\delta\rho_H)^2 \quad (23)$$

$$\leq (\delta\rho_L)^2 + \frac{1}{4^2}(\delta\rho_H)^2. \quad (24)$$

We assume fractional errors equal to the maximum errors in calibration at each site and arising from the systematic uncertainty in the detector response (8.5% in L1, and 4.5% for LHO). This results in a conservative estimate of the error at the largest observed SNR:

$$(\delta\rho)^2 \leq (0.085\rho_{\max})^2 + 0.0625(0.045\rho_{\max})^2. \quad (25)$$

The resulting error on N_G is

$$\delta N_G|_{\text{cal}} = 0.02 \text{ MWEG}. \quad (26)$$

Simulations of the contributions to δN_G from the random fluctuations of α , assuming the largest possible error in H1 (18%), made negligible contributions to the error due to calibration, as expected.

3. Uncertainties in the analysis pipeline

Since we use matched filtering to search for gravitational waves from inspiraling binaries, differences between the theoretical and the real waveforms could also adversely affect the results. These effects have been studied in great detail for binary neutron star systems [18–20]. The results indicate a $\sim 10\%$ loss of SNR due to inaccurate modeling of the waveforms for binaries in the mass range of interest. This feeds into our result through our measurement of the efficiency. We may be overestimating our sensitivity to real binary inspiral signals; the estimated loss of efficiency results in an error on N_G of

$$\delta N_G|_{\text{wave}} = +0/ - 0.02 \text{ MWEG}. \quad (27)$$

The effects of discreteness of the template placement, errors in the estimates of the power spectral density $S(f)$ used in the matched filter in Eq. (3), and trends in the instrumental noise are all accounted for by the Monte-

Carlo simulation. The error in the efficiency measurement due to the finite number of injections is estimated as

$$\delta N_G^i|_{\text{MC}} = L_i \sqrt{\varepsilon_i(1 - \varepsilon_i)/n^i} \quad (28)$$

where n^i is the number of injections made for each galaxy. The total error from the Monte Carlo is

$$\delta N_G|_{\text{MC}} = 0.02 \text{ MWEG}. \quad (29)$$

4. Combined uncertainties on N_G and the rate

Combining the errors in quadrature yields total errors

$$\delta N_G|_{\text{total}} = +0.06/ - 0.07 \text{ MWEG}. \quad (30)$$

To be conservative, we assume the downward excursion $N_G = 1.34 - 0.07 = 1.27$ MWEG when using Eq. (15) to derive an observational upper limit on the rate of binary neutron star coalescence

$$\mathcal{R} < 47 \text{ yr}^{-1} \text{ MWEG}^{-1} \quad (31)$$

where we have used $T = 339$ hr and $P_b = 1$.

X. SUMMARY AND DISCUSSION

Using data from the second LIGO science run (S2), we have significantly improved our methods and strategies to search for waveforms from inspiraling neutron stars. The search has been optimized for detecting signals, rather than for setting optimal upper limits on the rate of sources [5]. In order to increase confidence in a detection candidate, we used data only when two or more detectors are operating. We performed extensive validations of the detection efficiency of our search using software and hardware signal injections. Using these injections, we tuned the parameters in our search to take advantage of the increased reach of the detectors in the S2 run and have almost 100% efficiency for signals in the Milky Way and the magellanic clouds. (Because of the antenna pattern, we cannot be 100% efficient.) We also achieved an average 6% efficiency for sources in the Andromeda Galaxy, at a distance of 770 kpc, near the typical maximum range of our second best detector (H1, 900 kpc). Using the same detection pipeline, but shifting in time the data from the two observatories, we quantified the accidental rate for our pipeline. This allows evaluation of our confidence that any candidate is a true signal.

The search of 373 hours of double- and triple-detector data during S2 resulted in no detection: the strongest candidate from the search, arising from coincidence between L1 and H2, had a combined signal-to-noise of 9.4. Our background estimate suggests that noise has a 20% probability of causing a candidate with this signal-to-noise or greater. Moreover, the weight of evidence from follow-up investigations suggests that this trigger originated with a burst of noise which corresponded to a brief period of servo instability. The detailed investigations presented in

Sec. VIII exemplify typical follow-up of coincident triggers which appear unlikely to arise from background. We have found no evidence of a gravitational-wave event from binary neutron star inspiral. Without a detection, the 339 hours of nonplayground data were used to place an upper limit on the rate of binary neutron star coalescence in the Universe.

This search was sensitive to binary inspiral signals from neutron star coalescence in the Local Group and other galaxies at distances of up to about 1.5 Mpc. We used Monte-Carlo simulations to measure the efficiency of our search to nonspinning neutron stars in this population. We conclude that the rate of binary neutron star coalescence is $\mathcal{R} \leq 47 \text{ yr}^{-1} \text{ MWEg}^{-1}$ with 90% confidence. This rate limit is significantly greater than astrophysically plausible rates [6], but it illustrates the performance of the search. It should be noted that the analysis of the first science run [5] was optimized to yield the lowest upper limit, while the method presented here was not. In particular, the decision to only use data when the Livingston interferometer *and* one of the Hanford interferometers were operating limited the observation time to 1/3 of the total data collected. Analysis of the remaining data collected in coincidence with TAMA will be reported elsewhere [41].

In this paper, we have presented a data analysis strategy that could lead to a detection of gravitational waves from

binary neutron star inspirals. The methods used to validate the search illustrate the subtleties of the analysis of several detectors with different sensitivities and orientations. Moreover, the experience gained by following up the largest coincident triggers will be crucial input to investigations of event candidates that are identified in future searches.

ACKNOWLEDGMENTS

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

-
- [1] P.R. Saulson, *Fundamentals of Interferometric Gravitational Wave Detectors* (World Scientific, Singapore, 1994).
 - [2] B. C. Barish and R. Weiss, *Phys. Today* **52**(10), 44 (1999).
 - [3] B. Abbott *et al.* (LIGO Scientific Collaboration), *Nucl. Instrum. Methods, Sect. A* **517**, 154 (2004).
 - [4] C. Cutler *et al.*, *Phys. Rev. Lett.* **70**, 2984 (1993).
 - [5] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **69**, 122001 (2004).
 - [6] V. Kalogera *et al.*, *Astrophys. J.* **601**, L179 (2004).
 - [7] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **72**, 042002 (2005).
 - [8] M. Ando *et al.*, *Phys. Rev. Lett.* **86**, 3950 (2001).
 - [9] I. H. Stairs, *Science* **304**, 547 (2004).
 - [10] J. M. Weisberg and J. H. Taylor, in *Radio Pulsars*, ASP Conf. Series 302, edited by M. Bailes, D. J. Nice, and S. Thorsett (ASP, Provo, Utah, 2003).
 - [11] E. S. Phinney, *Astrophys. J.* **380**, L17 (1991).
 - [12] K. Belczynski, V. Kalogera, and T. Bulik, *Astrophys. J.* **572**, 407 (2002).
 - [13] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **72**, 082002 (2005).
 - [14] M. Mateo, *Annu. Rev. Astron. Astrophys.* **36**, 435 (1998).
 - [15] R. B. Tully, *VizieR Online Data Catalog* 7145, 0 (1994).
 - [16] G. de Vaucouleurs, A. de Vaucouleurs, H. G. Corwin, R. J. Buta, G. Paturel, and P. Fouque, *Third Reference Catalogue of Bright Galaxies (RC3)* (Springer-Verlag, New York, NY, 1991).
 - [17] P. Nutzman, V. Kalogera, L. S. Finn, C. Hendrickson, and K. Belczynski, *Astrophys. J.* **612**, 364 (2004).
 - [18] T. A. Apostolatos, *Phys. Rev. D* **52**, 605 (1995).
 - [19] S. Droz and E. Poisson, *Phys. Rev. D* **56**, 4449 (1997).
 - [20] S. Droz, *Phys. Rev. D* **59**, 064030 (1999).
 - [21] L. Bildsten and C. Cutler, *Astrophys. J.* **400**, 175 (1992).
 - [22] L. Blanchet, B. R. Iyer, C. M. Will, and A. G. Wiseman, *Classical Quantum Gravity* **13**, 575 (1996).
 - [23] S. Droz, D. J. Knapp, E. Poisson, and B. J. Owen, *Phys. Rev. D* **59**, 124016 (1999).
 - [24] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
 - [25] B. J. Owen and B. S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).
 - [26] B. Allen, *Phys. Rev. D* **71**, 062001 (2005).
 - [27] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown and J. D. E. Creighton, gr-qc/0509116.
 - [28] M. Ito, GlitchMon: A DMT monitor to look for transient signals in selected channels, developed using the LIGO Data Monitoring Tool (DMT) library.
 - [29] N. Christensen, P. Shawhan, and G. González, *Classical Quantum Gravity* **21**, S1747 (2004).
 - [30] LSC Algorithm Library software packages LAL and LALAPPS, the CVS tag versions iulgroup20040512 of LAL and LALAPPS were used in this analysis, <http://www.ligo.org>.

- isc-group.phys.uwm.edu/lal.
- [31] T. Tannenbaum, D. Wright, K. Miller, and M. Livny, in *Beowulf Cluster Computing with Linux*, edited by T. Sterling (MIT Press, Cambridge, MA, 2001).
 - [32] S. Fairhurst *et al.*, “Analysis of inspiral hardware injections during s2” (to be published).
 - [33] N. Christensen, R. Meyer, and A. Libson, *Classical Quantum Gravity* **21**, 317 (2004).
 - [34] A. Pai, S. Dhurandhar, and S. Bose, *Phys. Rev. D* **64**, 042004 (2001).
 - [35] P.R. Brady, J.D.E. Creighton, and A.G. Wiseman, *Classical Quantum Gravity* **21**, S1775 (2004).
 - [36] V. Kalogera, R. Narayan, D.N. Spergel, and J.H. Taylor, *Astrophys. J.* **556**, 340 (2001).
 - [37] S.E. Thorsett and D. Chakrabarty, *Astrophys. J.* **512**, 288 (1999).
 - [38] R. Adhikari, G. González, M. Landry, and B. O’Reilly, *Classical Quantum Gravity* **20**, S903 (2003).
 - [39] D.A. Brown (LIGO Scientific Collaboration), *Classical Quantum Gravity* **21**, S797 (2004).
 - [40] G. González, M. Landry, and B. O’Reilly, Report No. LIGO-T040060-00-D, LIGO Project, 2004, <http://www.ligo.caltech.edu/docs/T/T040060-00.pdf>.
 - [41] B. Abbott *et al.* (LIGO-TAMA) (to be published).