First all-sky upper limits from LIGO on the strength of periodic gravitational waves using the Hough transform

B. Abbott,[12] R. Abbott,[15] R. Adhikari,[12] A. Ageev,[20,27] J. Agresti,[12] B. Allen,[39] J. Allen,[13] R. Amin,[16] 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,[15] R. Bork,[12] S. Bose,[41] 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] J. Cannizzo,[21] K. Cannon,[39] 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,[12] S. Chelkowski,[2,31] Y. Chen,[1] 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] J. Dalrymple,[27] E. D'Ambrosio,[12] K. Danzmann,[31,2] G. Davies,[7] E. Daw,[16,k] D. DeBra,[26] T. Delker,[34,l] V. Dergachev,[36] S. Desai,[28] R. DeSalvo,[12] S. Dhurandhar,[11] A. Di Credico,[27] M. Díaz,[29] H. Ding,[12] R. W. P. Drever,[4] R. J. Dupuis,[12] 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] L. Goggin,[12] G. González,[16] S. Goßler,[31] P. Grandclément,[23,o] A. Grant,[35] C. Gray,[14] A. M. Gretarsson,[15,p] D. Grimmett,[12] H. Grote,[2] S. Grunewald,[1] M. Guenther,[14] E. Gustafson,[26,q] 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,r] W. Hua,[26] M. Ito,[37] Y. Itoh,[1] A. Ivanov,[12] O. Jennrich,[35,s] B. Johnson,[14] W. W. Johnson,[16] W. R. Johnston,[29] D. I. Jones,[28] G. Jones,[7] L. Jones,[12] D. Jungwirth,[12,t] V. Kalogera,[23] E. Katsavounidis,[13] K. Kawabe,[14] S. Kawamura,[22] W. Kells,[12] J. Kern,[15,u] A. Khan,[15] S. Killbourn,[35] C. J. Killow,[35] C. Kim,[23] C. King,[12] P. King,[12] S. Klimenko,[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,v] M. Lormand,[15] M. Lubinski,[14] H. Lück,[31,2] M. Luna,[32] T. T. Lyons,[12,v] B. Machenschalk,[1] M. MacInnis,[13] M. Mageswaran,[12] K. Mailand,[12] W. Majid,[12,b] M. Malec,[2,31] V. Mandic,[12] F. Mann,[12] A. Marin,[13,w] S. Márka,[9,x] E. Maros,[12] J. Mason,[12,y] K. Mason,[13] O. Matherny,[14] L. Matone,[9] N. Mavalvala,[13] R. McCarthy,[14] D. E. McClelland,[3] M. McHugh,[18] J. W. C. McNabb,[28] A. Melissinos,[38] G. Mendell,[14] R. A. Mercer,[33] S. Meshkov,[12] E. Messaritaki,[39] C. Messenger,[33] E. Mikhailov,[13] S. Mitra,[11] V. P. Mitrofanov,[20] G. Mitselmakher,[34] R. Mittleman,[13] O. Miyakawa,[12] S. Miyoki,[12,z] S. Mohanty,[29] G. Moreno,[14] K. Mossavi,[2] G. Mueller,[34] S. Mukherjee,[29] P. Murray,[35] E. Myers,[40] J. Myers,[14] S. Nagano,[2] T. Nash,[12] R. Nayak,[11] G. Newton,[35] F. Nocera,[12] J. S. Noel,[41] P. Nutzman,[23] T. Olson,[24] B. O'Reilly,[15] D. J. Ottaway,[13] A. Ottewill,[39,aa] D. Ouimette,[12,t] H. Overmier,[15] B. J. Owen,[28] Y. Pan,[6] M. A. Papa,[1] V. Parameshwaraiah,[14] A. Parameswaran,[2] 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,ab] R. Riesen,[15] K. Riles,[36] B. Rivera,[14] A. Rizzi,[15,ac] D. I. Robertson,[35] N. A. Robertson,[26,35] C. Robinson,[7] L. Robison,[12] S. Roddy,[15] A. Rodriguez,[16] J. Rollins,[9] J. D. Romano,[7] J. Romie,[12] H. Rong,[34,n] D. Rose,[12] E. Rotthoff,[28] S. Rowan,[35] A. Rüdiger,[2] L. Ruet,[13] P. Russell,[12] K. Ryan,[14] I. Salzman,[12] V. Sandberg,[14] G. H. Sanders,[12,ad] V. Sannibale,[12] P. Sarin,[13] B. Sathyaprakash,[7] P. R. Saulson,[27] R. Savage,[14] A. Sazonov,[34] R. Schilling,[2] K. Schlaufman,[28] V. Schmidt,[12,ae] R. Schnabel,[19] R. Schofield,[37] B. F. Schutz,[1,7] P. Schwinberg,[14] S. M. Scott,[3] S. E. Seader,[41] A. C. Searle,[3] B. Sears,[12] S. Seel,[12] F. Seifert,[19] D. Sellers,[15] A. S. Sengupta,[11] C. A. Shapiro,[28,af] P. Shawhan,[12] D. H. Shoemaker,[13] Q. Z. Shu,[34,ag] 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] O. Spjeld,[15] G. Stapfer,[15] D. Steussy,[8] K. A. Strain,[35] D. Strom,[37] A. Stuver,[28] T. Summerscales,[28] M. C. Sumner,[12] M. Sung,[16] P. J. Sutton,[12] J. Sylvestre,[12,ah] A. Takamori,[12,ai] 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,aj] 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,ak] 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] R. Ward,[12] B. Ware,[12,b] K. Watts,[15]

D. Webber,[12] A. Weidner,[19,2] U. Weiland,[31] A. Weinstein,[12] R. Weiss,[13] H. Welling,[31] L. Wen,[1] S. Wen,[16] K. Wette,[3] 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,al] R. Williams,[4] B. Willke,[31,2] A. Wilson,[12] B. J. Winjum,[28,e] W. Winkler,[2] S. Wise,[34] A. G. Wiseman,[39] G. Woan,[35] D. Woods,[39] 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,am] L. Zhang,[12] R. Zhu,[1] N. Zotov,[17] M. Zucker,[15] and J. Zweizig[12]

(LIGO Scientific Collaboration)[an]

[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]*Columbia University, New York, New York 10027, 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]*Vassar College, Poughkeepsie, New York 12604, USA*
[41]*Washington State University, Pullman, Washington 99164, USA*

(Received 13 September 2005; published 28 November 2005)

We perform a wide parameter-space search for continuous gravitational waves over the whole sky and over a large range of values of the frequency and the first spin-down parameter. Our search method is based on the Hough transform, which is a semicoherent, computationally efficient, and robust pattern recognition technique. We apply this technique to data from the second science run of the LIGO detectors and our final results are all-sky upper limits on the strength of gravitational waves emitted by unknown isolated spinning neutron stars on a set of narrow frequency bands in the range 200–400 Hz. The best upper limit on the gravitational-wave strain amplitude that we obtain in this frequency range is 4.43×10^{-23}.

DOI: 10.1103/PhysRevD.72.102004 PACS numbers: 04.80.Nn, 07.05.Kf, 95.55.Ym, 97.60.Gb

I. INTRODUCTION

Continuous gravitational signals emitted by rotating neutron stars are promising sources for interferometric gravitational-wave detectors such as GEO 600 [1,2], the Laser Interferometer Gravitational-Wave Observatory (LIGO) [3,4], TAMA 300 [5] and VIRGO [6]. There are several physical mechanisms which might cause a neutron star to emit periodic gravitational waves. The main possibilities considered in the literature are (i) nonaxisymmetric distortions of the solid part of the star [7–10], (ii) unstable r modes in the fluid [7,11,12], and (iii) free precession (or ''wobble'') [13,14]. The detectability of a signal depends on the detector sensitivity, the intrinsic emission strength, the source distance and its orientation. If the source is not known, the detectability also depends on the available computational resources. For some search methods the detectability of a signal also depends on the source model used, but an all-sky wide-band search such as detailed here can detect any of the sources described above.

Previous searches for gravitational waves from rotating neutron stars have been of two kinds. The first is a search targeting pulsars whose parameters are known through radio observations. These searches typically use matched filtering techniques and are not very computationally expensive. Examples of such searches are [15,16] which targeted known radio pulsars, at twice the pulsar frequency, using data from the first and second science runs of the GEO 600 and LIGO detectors [17]. No signals were detected and the end results were upper limits on the strength of the gravitational waves emitted by these pulsars and therefore on their ellipticity.

The second kind of search looks for as yet undiscovered rotating neutron stars. An example of such a search is [18] in which a two-day long data stretch from the Explorer bar detector is used to perform an all-sky search in a narrow frequency band around the resonant frequency of the detector. Another example is [19] which uses data from the LIGO detectors to perform an all-sky search in a wide frequency band using 10 h of data. The same paper also describes a search for a gravitational-wave signal from the compact companion to Sco X-1 in a large *orbital* parameter space using 6 h of data. The key issue in these wide parameter-space searches is that a fully coherent all-sky search over a large frequency band using a significant amount of data is computationally limited. This is because looking for weak continuous wave signals requires long observation times to build up sufficient signal-to-noise ratio; the amplitude signal-to-noise ratio increases as the square root of the observation time. On the other hand, the number of templates that must be considered, and therefore the computational requirements, scale much faster than linearly with the observation time. We therefore need methods which are suboptimal but computationally less expensive [20–24]. Such methods typically involve semi-coherent combinations of the signal power in short stretches of data, and the Hough transform is an example of such a method.

The Hough transform is a pattern recognition algorithm which was originally invented to analyze bubble chamber pictures from CERN [25]. It was later patented by IBM [26], and it has found many applications in the analysis of digital images [27]. A detailed discussion of the Hough transform as applied to the search for continuous gravitational waves can be found in [23,28]. In this paper, we apply this technique to perform an all-sky search for isolated spinning neutron stars using two months of data collected in early 2003 from the second science run of the LIGO detectors (henceforth denoted as the S2 run). The

[a]Present address: Stanford Linear Accelerator Center.
[b]Present address: Jet Propulsion Laboratory.
[c]Permanent address: HP Laboratories.
[d]Present address: Rutherford Appleton Laboratory.
[e]Present address: University of California, Los Angeles.
[f]Present address: Hofstra University.
[g]Permanent address: GReCO, Institut d'Astrophysique de Paris (CNRS).
[h]Present address: Charles Sturt University, Australia.
[i]Present address: Keck Graduate Institute.
[j]Present address: National Science Foundation.
[k]Present address: University of Sheffield.
[l]Present address: Ball Aerospace Corporation.
[m]Present address: European Gravitational Observatory.
[n]Present address: Intel Corp..
[o]Present address: University of Tours, France.
[p]Present address: Embry-Riddle Aeronautical University.
[q]Present address: Lightconnect Inc..
[r]Present address: W.M. Keck Observatory.
[s]Present address: ESA Science and Technology Center.
[t]Present address: Raytheon Corporation.
[u]Present address: New Mexico Institute of Mining and Technology/Magdalena Ridge Observatory Interferometer.
[v]Present address: Mission Research Corporation.
[w]Present address: Harvard University.
[x]Permanent address: Columbia University.
[y]Present address: Lockheed-Martin Corporation.
[z]Permanent address: University of Tokyo, Institute for Cosmic Ray Research.
[aa]Permanent address: University College Dublin.
[ab]Present address: Research Electro-Optics Inc..
[ac]Present address: Institute of Advanced Physics, Baton Rouge, LA.
[ad]Present address: Thirty Meter Telescope Project at Caltech.
[ae]Present address: European Commission, DG Research, Brussels, Belgium.
[af]Present address: University of Chicago.
[ag]Present address: LightBit Corporation.
[ah]Permanent address: IBM Canada Ltd..
[ai]Present address: The University of Tokyo.
[aj]Present address: University of Delaware.
[ak]Permanent address: Jet Propulsion Laboratory.
[al]Present address: Shanghai Astronomical Observatory.
[am]Present address: Laser Zentrum Hannover.
[an]Electronic address: http://www.ligo.org.

main results of this paper are all-sky upper limits on a set of narrow frequency bands within the range 200–400 Hz and including one spin-down parameter.

Our best 95% frequentist upper limit on the gravitational-wave strain amplitude for this frequency range is 4.43×10^{-23}. As discussed below in Sec. III, based on the statistics of the neutron star population with optimistic assumptions, this upper limit is about 1 order of magnitude larger than the amplitude of the strongest expected signal. For this reason, the present search is unlikely to discover any neutron stars, and we focus here on setting upper limits. As we shall see later, with 1 yr of data at design sensitivity for initial LIGO, we should gain about 1 order of magnitude in sensitivity, thus enabling us to detect signals smaller than what is predicted by the statistical argument mentioned above. Substantial improvements in the detector noise have been achieved since the S2 observations. A third science run (S3) took place at the end of 2003 and a fourth science run (S4) at the beginning of 2005. In these later runs LIGO instruments collected data of improved sensitivity, but still less sensitive than the instruments' design goal. Several searches for various types of gravitational-wave signals have been completed or are under way using data from the S2 and S3 runs [16,19,29–35]. We expect that the method presented here, applied as part of a hierarchical scheme and used on a much larger data set, will eventually enable the direct detection of periodic gravitational waves.

This paper is organized as follows: Sec. II describes the second science run of the LIGO detectors; Sec. III summarizes the current understanding of the astrophysical targets; Sec. IV reviews the waveform from an isolated spinning neutron star; Sec. V presents the general idea of our search method, the Hough transform, and summarizes its statistical properties; Sec. VI describes its implementation and results on short Fourier transformed data. The upper limits are given in Sec. VII. Section VIII presents a validation of our search method using hardware injected signals, and finally Sec. IX concludes with a summary of our results and suggestions for further work.

II. THE SECOND SCIENCE RUN

The LIGO detector network consists of a 4 km interferometer in Livingston, Louisiana (L1), and two interferometers in Hanford, Washington, one 4 km and the other 2 km (H1 and H2). Each detector is a power-recycled Michelson interferometer with long Fabry-Perot cavities in each of its orthogonal arms. These interferometers are sensitive to quadrupolar oscillations in the space-time metric due to a passing gravitational wave, measuring directly the gravitational-wave strain amplitude.

The data analyzed in this paper were produced during LIGO's 59 d second science run. This run started on February 14 and ended April 14, 2003. Although the GEO detector was not operating at the time, all three LIGO detectors were functioning at a significantly better sensitivity than during S1, the first science run [17], and had displacement spectral amplitudes near 10^{-18} m-Hz$^{-1/2}$ between 200 and 400 Hz. The strain sensitivities in this science run were within an order of magnitude of the design sensitivity for the LIGO detectors. For a description of the detector configurations for S2 we refer the reader to [29] Sec. IV and [30] Sec. II.

The reconstruction of the strain signal from the error signal of the feedback loop, used to control the differential length of the interferometer arms, is referred to as the calibration. Changes in the calibration were tracked by injecting continuous, fixed-amplitude sinusoidal excitations into the end test mass control systems, and monitoring the amplitude of these signals at the measurement error point. Calibration uncertainties at the three LIGO detectors during S2 were estimated to be smaller than 11% [36].

The data were acquired and digitized at a rate of 16 384 Hz. The duty cycle for the interferometers, defined as the fraction of the total run time when the interferometer was locked (i.e., all interferometer control servos operating in their linear regime) and in its low noise configuration, were similar to those of the previous science run, approximately 37% for L1, 74% for H1 and 58% for H2. The longest continuous locked stretch for any interferometer during S2 was 66 h for H1. The smaller duty cycle for L1 was due to anthropogenic diurnal low-frequency seismic noise which prevented operations during the day on weekdays. Recently installed active feedback seismic isolation has successfully addressed this problem.

Figure 1 shows the expected sensitivity for the Hough search by the three LIGO detectors during S2. Those h_0 values correspond to the amplitudes detectable from a generic continuous gravitational-wave source, if we were performing a targeted search, with a 1% false alarm rate and 10% false dismissal rate, as given by Eq. (17). The differences among the three interferometers reflect differences in the operating parameters, hardware implementation of the three instruments, and the duty cycles. Figure 1 also shows the expected sensitivity (at the same false alarm and false dismissal rates) for initial LIGO 4 km interferometers running at design sensitivity assuming an observation time of 1 yr. These false alarm and false dismissal values are chosen in agreement with [15,19] only for comparison purposes. Because of the large parameter search we perform here, it would be more meaningful to consider a lower false alarm rate, say of 10^{-10} and then the sensitivity for a targeted search would get worse by a factor 1.5. The search described in this paper is not targeted and this degrades the sensitivity even further. This will be discussed in Sec. VII.

At the end of the S2 run, two fake artificial pulsar signals were injected for a 12 h period into all three LIGO interferometers. These hardware injections were done by modulating the mirror positions via the actuation control

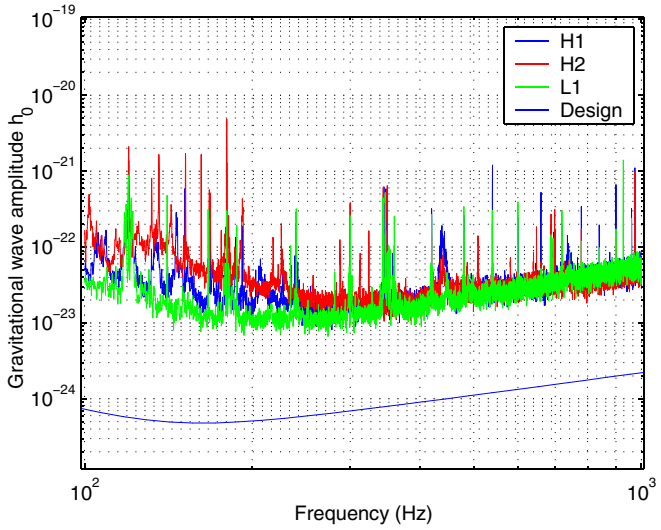

FIG. 1 (color online). Characteristic amplitude detectable from a known generic source with a 1% false alarm rate and 10% false dismissal rate, as given by Eq. (17). All curves use typical sensitivities of the three LIGO detectors during S2 and observation times corresponding to the up-time of the detectors during S2. The thin line is the expected characteristic amplitude for the same false alarm and false dismissal rates, but using the initial LIGO design goal for the 4 km instruments and an effective observation time of 1 yr.

signals. These injections were designed to give an end-to-end validations of the search pipelines starting from as far up the observing chain as possible. See Sec. VIII for details.

III. ASTROPHYSICAL TARGETS

The target population of this search consists of isolated rotating neutron stars that are not observed in electromagnetic waves. Current models of stellar evolution suggest that our Galaxy contains of order 10^9 neutron stars and that of order 10^5 are active pulsars [37]. Up to now, only of order 10^3 objects have been identified as neutron stars, either by observation as pulsars, or through their x-ray emission, of which about 90% are isolated [38–41]. Most neutron stars will remain unobserved electromagnetically for many reasons such as the nonpulsed emission being faint or the pulses being emitted in a beam which does not sweep across the Earth. Therefore, there are many more neutron stars in the target population than have already been observed.

Although there is great uncertainty in the physics of the emission mechanism and the strength of an individual source, we can argue for a robust upper limit on the strength of the strongest source in the galactic population that is almost independent of individual source physics. The details of the argument and an overview of emission mechanisms can be found in a forthcoming paper [19]. Here we do not repeat the details but merely summarize the result. For an upper limit we make optimistic assumptions—that neutron stars are born rapidly rotating and spinning down due to gravitational waves, and that they are distributed uniformly throughout the galactic disc—and the plausible assumption that the overall galactic birthrate $1/\tau_b$ is steady. By converting these assumptions to a distribution of neutron stars with respect to gravitational-wave strain and frequency, we find there is a 50% chance that the strongest signal between frequencies $f_{\min}$ and $f_{\max}$ has an amplitude of at least

$$h_0 \approx 4 \times 10^{-24} \left[\left(\frac{30\ \mathrm{yr}}{\tau_b} \right) \ln \frac{f_{\max}}{f_{\min}} \right]^{1/2}. \quad (1)$$

Of course, with less optimistic assumptions this value would be smaller.

Comparing Eq. (1) to Fig. 1, a search of S2 data is not expected to result in a discovery. However, it is still possible that the closest neutron star is closer than the typical distance expected from a random distribution of supernovae (for example due to recent star formation in the Gould belt as considered in Ref. [42]). It is also possible that a "blind" search of this sort may discover some previously unknown class of compact objects not born in supernovae. More importantly, future searches for previously undiscovered rotating neutron stars using the methods presented here will be much more sensitive. The goal of initial LIGO is to take a year of data at design sensitivity, which means a factor 10 decrease in the amplitude strain noise relative to S2, and a factor 10 increase in the length of the data set. These combine to reduce h_0 to somewhat below the value in Eq. (1), and thus initial LIGO at full sensitivity will have some chance of observing a periodic signal.

IV. THE EXPECTED WAVEFORM

In order to describe the expected signal waveform we will use the same notation as [15]. We will briefly summarize it in the next paragraphs for convenience. The form of the gravitational wave emitted by an isolated spinning neutron star, as seen by a gravitational-wave detector, is

$$h(t) = F_+(t, \psi) h_+(t) + F_\times(t, \psi) h_\times(t), \quad (2)$$

where t is time in the detector frame, ψ is the polarization angle of the wave, and $F_{+,\times}$ are the detector antenna pattern functions for the two polarizations. If we assume the emission mechanism is due to deviations of the pulsar's shape from perfect axial symmetry, then the gravitational waves are emitted at a frequency which is exactly twice the rotational rate f_r. Under this assumption, the waveforms for the two polarizations $h_{+,\times}$ are given by

$$h_+ = h_0 \frac{1 + \cos^2 \iota}{2} \cos\Phi(t), \quad (3)$$

$$h_\times = h_0 \cos\iota \sin\Phi(t), \quad (4)$$

where ι is the angle between the neutron star's spin axis and the direction of propagation of the waves, and h_0 is the amplitude:

$$h_0 = \frac{16\pi^2 G}{c^4}\frac{I_{zz}\epsilon f_r^2}{d}, \tag{5}$$

where G is Newton's gravitational constant, c the speed of light, I_{zz} is the principal moment with the z axis being its spin axis, $\epsilon := (I_{xx} - I_{yy})/I_{zz}$ is the equatorial ellipticity of the star, and d is the distance to the star.

The phase $\Phi(t)$ takes its simplest form in the Solar System Barycenter (SSB) frame where it can be expanded in a Taylor series up to second order:

$$\Phi(t) = \Phi_0 + 2\pi(f_0(T - T_0) + \tfrac{1}{2}\dot{f}(T - T_0)^2). \tag{6}$$

Here T is time in the SSB frame and T_0 is a fiducial start time. The phase Φ_0, frequency f_0 and spin-down parameter $\dot{f}$ are defined at this fiducial start time. In this paper, we include only one spin-down parameter in our search; as we shall see later in Sec. VI B, our frequency resolution is too coarse for the higher spin-down parameters to have any significant effect on the frequency evolution of the signal (for the spin-down ages we consider).

Modulo relativistic effects which are unimportant for this search, the relation between the time of arrival T of the wave in the SSB frame and in the detector frame t is

$$T = t + \frac{\mathbf{n}\cdot\mathbf{r}}{c}, \tag{7}$$

where $\mathbf{n}$ is the unit vector from the detector to the neutron star, and $\mathbf{r}$ is the detector position in the SSB frame.

The instantaneous frequency $f(t)$ of the wave as observed by the detector is given, to a very good approximation, by the familiar nonrelativistic Doppler formula:

$$f(t) - \hat{f}(t) = \hat{f}(t)\frac{\mathbf{v}(t)\cdot\mathbf{n}}{c}, \tag{8}$$

where $\hat{f}(t)$ is the instantaneous signal frequency in the SSB frame at time t:

$$\hat{f}(t) = f_0 + \dot{f}\left(t - t_0 + \frac{\Delta\mathbf{r}(t)\cdot\mathbf{n}}{c}\right), \tag{9}$$

where t_0 is the fiducial detector time at the start of the observation and $\Delta\mathbf{r}(t) = \mathbf{r}(t) - \mathbf{r}(t_0)$. It is easy to see that the $\Delta\mathbf{r}\cdot\mathbf{n}/c$ term can safely be ignored so that, to an excellent approximation

$$\hat{f}(t) = f_0 + \dot{f}(t - t_0). \tag{10}$$

V. THE HOUGH TRANSFORM

In this paper, we use the Hough transform to find the pattern produced by the Doppler shift (8) and the spin down (10) of a gravitational-wave signal in the time-frequency plane of our data. This pattern is independent of the source model used and therefore of the emission mechanisms. We only assume that the gravitational-wave signal is emitted by an isolated spinning neutron star.

The starting point for our search is a set of data segments, each corresponding to a time interval T_{coh}. Each of these data segments is Fourier transformed to produce a set of N short time-baseline Fourier transforms (SFTs). From this set of SFTs, calculating the periodograms (the square modulus of the Fourier transform) and selecting frequency bins (peaks) above a certain threshold, we obtain a time-frequency map of our data. In the absence of a signal the peaks in the time-frequency plane are distributed in a random way; if signal is present, with high enough signal-to-noise ratio, some of these peaks will be distributed along the trajectory of the received frequency of the signal.

The Hough transform maps points of the time-frequency plane into the space of the source parameters $(f_0, \dot{f}, \mathbf{n})$. The result of the Hough transform is a histogram, i.e., a collection of integer numbers, each representing the detection statistic for each point in parameter space. We shall refer to these integers as the number count. The number counts are computed in the following way: For each selected bin in the SFTs, we find which points in parameter space are consistent with it, according to Eq. (8), and the number count in all such points is increased by unity. This is repeated for all the selected bins in all the SFTs to obtain the final histogram.

To illustrate this, let us assume the source parameters are only the coordinates of the source in the sky, and this source is emitting a signal at a frequency f_0. Moreover we assume that at a given time t a peak at frequency f has been selected in the corresponding SFT. The Hough transform maps this peak into the loci of points, on the celestial sphere, where a source emitting a signal with frequency f_0 could be located in order in order to produce at the detector a peak at f. By repeating this for all the selected peaks in our data we will obtain the final Hough map. If the peaks in the time-frequency plane were due only to signal, all the corresponding loci would intersect in a region of the Hough map identifying the source position.

An advantage of the Hough transform is that a large region in parameter space can be analyzed in a single pass. By dropping the amplitude information of the selected peaks, the Hough search is expected to be computationally efficient, but at the cost of being somewhat less sensitive than others semicoherent methods, e.g., the stack-slide search [21]. On the other hand, discarding this extra information makes the Hough transform more robust against transient spectral disturbances because no matter how large a spectral disturbance is in a single SFT, it will contribute at the most $+1$ to the number count. This is not surprising since the optimal statistic for the detection of weak signals in the presence of a Gaussian background with large non Gaussian outliers is effectively cut off above some value

[43,44]. This is, in practice, what the Hough transform does to large spectral outliers.

With the above short summary at hand, we now give the relevant notation and equations that will be used later. For further details and derivations of the equations below, we refer the reader to [23].

Frequency bins are selected by setting a threshold $\rho_{\rm th}$ on the normalized power ρ_k defined as

$$\rho_k = \frac{2|\tilde{x}_k|^2}{T_{\rm coh} S_n(f_k)}, \tag{11}$$

where $\tilde{x}_k$ is the discrete Fourier transform of the data, the frequency index k corresponds to a physical frequency of $f_k = k/T_{\rm coh}$, and $S_n(f_k)$ is the single sided power spectral density of the detector noise. The kth frequency bin is selected if $\rho_k \geq \rho_{\rm th}$, and rejected otherwise. In this way, each SFT is replaced by a collection of zeros and ones called a peak-gram.

Let n be the number count at a point in parameter space, obtained after performing the Hough transform on our data. Let $p(n)$ be the probability distribution of n in the absence of a signal, and $p(n|h)$ the distribution in the presence of a signal $h(t)$. It is clear that $0 \leq n \leq N$, and it can be shown that for stationary Gaussian noise, $p(n)$ is a binomial distribution with mean Nq where q is the probability that any frequency bin is selected:

$$p(n) = \binom{N}{n} q^n (1-q)^{N-n}. \tag{12}$$

For Gaussian noise in the absence of a signal, it is easy to show that ρ_k follows an exponential distribution so that $q = e^{-\rho_{\rm th}}$. In the presence of a signal, the distribution is ideally also a binomial but with a slightly larger mean $N\eta$ where, for weak signals, η is given by

$$\eta = q\left\{1 + \frac{\rho_{\rm th}}{2}\lambda + \mathcal{O}(\lambda^2)\right\}. \tag{13}$$

λ is the signal-to-noise ratio within a single SFT, and for the case when there is no mismatch between the signal and the template:

$$\lambda = \frac{4|\tilde{h}(f_k)|^2}{T_{\rm coh} S_n(f_k)}, \tag{14}$$

with $\tilde{h}(f)$ being the Fourier transform of the signal $h(t)$. The approximation that the distribution in the presence of a signal is binomial breaks down for reasonably strong signals. This is due to possible nonstationarities in the noise, and the amplitude modulation of the signal which causes η to vary from one SFT to another.

Candidates in parameter space are selected by setting a threshold $n_{\rm th}$ on the number count. The false alarm and false dismissal rates for this threshold are defined, respectively, in the usual way:

$$\alpha = \sum_{n=n_{\rm th}}^{N} p(n), \qquad \beta = \sum_{n=0}^{n_{\rm th}-1} p(n|h). \tag{15}$$

We choose the thresholds $(n_{\rm th}, \rho_{\rm th})$ based on the Neyman-Pearson criterion of minimizing β for a given value of α. It can be shown [23] that this criteria leads, in the case of weak signals, large N, and Gaussian stationary noise, to $\rho_{\rm th} \approx 1.6$. This corresponds to $q \approx 0.20$, i.e., we select about 20% of the frequency bins from each SFT. This value of $\rho_{\rm th}$ turns out to be independent of the choice of α and signal strength. Furthermore, $n_{\rm th}$ is also independent of the signal strength and is given by

$$n_{\rm th} = Nq + \sqrt{2Nq(1-q)}\,{\rm erfc}^{-1}(2\alpha), \tag{16}$$

where $\rm erfc^{-1}$ is the inverse of the complementary error function. These values of the thresholds lead to a false dismissal rate β which is given in [23]. The value of β of course depends on the signal strength, and on the average, the weakest signal which will cross the above thresholds at a false alarm rate α and false dismissal β is given by

$$h_0 = 5.34 \frac{S^{1/2}}{N^{1/4}} \sqrt{\frac{S_n}{T_{\rm coh}}}, \tag{17}$$

where

$$S = {\rm erfc}^{-1}(2\alpha) + {\rm erfc}^{-1}(2\beta). \tag{18}$$

Equation (17) gives the smallest signal which can be detected by the search, and is therefore a measure of the sensitivity of the search.

VI. THE SEARCH

A. The SFT data

The input data to our search is a collection of calibrated SFTs with a time baseline $T_{\rm coh}$ of 30 min. While a larger value of $T_{\rm coh}$ leads to better sensitivity, this time baseline cannot be made arbitrarily large because of the frequency drift caused by the Doppler effect (and also the spin down); we would like the signal power of a putative signal to be concentrated in less than half the frequency resolution $1/T_{\rm coh}$. It is shown in [23] that at 300 Hz, we could ideally choose $T_{\rm coh}$ up to $\sim$60 min. On the other hand, we should be able to find a significant number of such data stretches during which the interferometers are in lock, the noise is stationary, and the data are labeled satisfactory according to certain data quality requirements. Given the duty cycles of the interferometers during S2 and the nonstationarity of the noise floor, it turns out that $T_{\rm coh} = 30$ min is a good compromise which satisfies these constraints. By demanding the data in each 30 min stretch to be continuous (although there could be gaps in between the SFTs) the number N of SFTs available for L1 data is 687, 1761 for H1 and 1384 for H2, reducing the nominal duty cycle for this search.

The SFT data are calibrated in the frequency domain by constructing a response function $R(f, t)$ that acts on the error signal of the feedback loop used to control the differential length of the interferometer arms. The response function $R(f, t)$ varies in time, primarily due to changes in the amount of light stored in the Fabry-Perot cavities of the interferometers. During S2, changes in the response were computed every 60 sec, and an averaging procedure was used to estimate the response function used on each SFT. The SFTs are windowed and high-pass filtered as described in Sec. IV C 1 of [15]. No further data conditioning is applied, although the data are known to contain many spectral disturbances, including the 60 Hz power line harmonics and the thermally excited violin modes of test mass suspension wires.

B. The parameter space

This section describes the portion of parameter space $(f_0, \dot{f}, \mathbf{n})$ we search over, and the resolution of our grid. Our template grid is not based on a metric calculation (as in e.g. [20,21]), but rather on a cubic grid which covers the parameter space as described below. Particular features of this grid are used to increase computational efficiency as described in Sec. VI C.

We analyze the full data set from the S2 run with a total observation time $T_{\rm obs} \sim 5.1 \times 10^6$ sec . The exact value of $T_{\rm obs}$ is different for the three LIGO interferometers [45]. We search for isolated neutron star signals in the frequency range 200–400 Hz with a frequency resolution

$$\delta f = \frac{1}{T_{\rm coh}} = 5.556 \times 10^{-4}\ \text{Hz}. \tag{19}$$

The choice of the range 200–400 Hz for the analysis is motivated by the low noise level, and therefore our ability to set the best upper limits for h_0, as seen from Fig. 1.

The resolution $\delta\dot{f}$ in the space of first spin-down parameters is given by the smallest value of $\dot{f}$ for which the intrinsic signal frequency does not drift by more than a single frequency bin during the total observation time [46]:

$$\delta\dot{f} = \frac{\delta f}{T_{\rm obs}} = \frac{1}{T_{\rm obs} T_{\rm coh}} \sim 1.1 \times 10^{-10}\ \text{Hz-s}^{-1}. \tag{20}$$

We choose the range of values $-\dot{f}_{\rm max} \le \dot{f} \le 0$, where the largest spin-down parameter $\dot{f}_{\rm max}$ is about 1.1×10^{-9}Hz-s^{-1}. This yields 11 spin-down values for each intrinsic frequency. In other words, we look for neutron stars whose spin-down age is at least $\tau_{\rm min} = \hat{f}/\dot{f}_{\rm max}$. This corresponds to a minimum spin-down age of 5.75×10^3 yr at 200 Hz, and 1.15×10^4 yr at 400 Hz. These values of $\dot{f}_{\rm max}$ and $\tau_{\rm min}$ are such that all known pulsars have a smaller spin-down rate than $\dot{f}_{\rm max}$ and, except for a few supernova remnants, all of them have a spin-down age significantly greater than the numbers quoted above. With these values of $\tau_{\rm min}$, it is easy to see that the second spin-down parameter can be safely neglected; it would take about 10 yr for the largest second spin-down parameter to cause a frequency drift of half a frequency bin.

As described in [23], for every given time, value of the intrinsic frequency f_0 and spin down $\dot{f}$, the set of sky locations $\mathbf{n}$ consistent with a selected frequency $f(t)$ corresponds to a constant value of $\mathbf{v} \cdot \mathbf{n}$ given by (8). This is a circle in the celestial sphere. It can be shown that every frequency bin of width δf corresponds to an annulus on the celestial sphere whose width is at least

$$(\delta\theta)_{\rm min} = \frac{c}{v} \frac{\delta f}{\hat{f}}, \tag{21}$$

with v being the magnitude of the average velocity of the detector in the SSB frame.

The resolution $\delta\theta$ in sky positions is chosen to be frequency dependent, being at most $\delta\theta = \frac{1}{2}(\delta\theta)_{\rm min}$. To choose the template spacing only, we use a constant value of v/c equal to 1.06×10^{-4}. This yields:

$$\delta\theta = 9.3 \times 10^{-3}\text{rad} \times \left(\frac{300\ \text{Hz}}{\hat{f}}\right). \tag{22}$$

This resolution corresponds to approximately 1.5×10^5 sky locations for the whole sky at 300 Hz. For that, we break up the sky into 23 sky patches of roughly equal area and, by means of the stereographic projection, we map each portion to a plane, and set a uniform grid with spacing $\delta\theta$ in this stereographic plane. The stereographic projection maps circles in the celestial sphere to circles in the plane thereby mapping the annuli in the celestial sphere, described earlier, to annuli in the stereographic plane. We ensure that the dimensions of each sky-patch are sufficiently small so that the distortions produced by the stereographic projection are not significant. This is important to ensure that the number of points needed to cover the full sky is not much larger than if we were using exactly the frequency resolution given by Eq. (22).

This adds up to a total number of templates per 1 Hz band at 200 Hz $\sim 1.9 \times 10^9$ while it increases up to 7.5×10^9 at 400 Hz.

C. The implementation of the Hough transform

This section describes in more detail the implementation of the search pipeline which was summarized in Sec. V. The first step in this semicoherent Hough search is to select frequency bins from the SFTs and construct the peak-grams. As mentioned in Sec. V, our criteria for selecting frequency bins is to set a threshold of 1.6 on the normalized power (11), thereby selecting about 20% of the frequency bins in every SFT.

The power spectral density S_n appearing in Eq. (11) is estimated by means of a running median applied to the periodogram of each individual SFT. The window size we employ for the running median is $w = 101$ corresponding

to 0.056 Hz [47]. The running median is a robust method to estimate the noise floor [48–50] which has the virtue of discarding outliers which appear in a small number of bins, thereby providing an accurate estimate of the noise floor in the presence of spectral disturbances and possible signals. The use of the median (instead of the mean) to estimate the power spectral density introduces a minor technical complication (see Appendix A for further details).

The next step is to choose a tiling of the sky. As described before, we break up the sky into 23 patches, of roughly equal area. By means of the stereographic projection, we map each portion to a two dimension plane and set a uniform grid with a resolution $\delta\theta$ in this plane. All of our calculations are performed on this stereographic plane, and are finally projected back on to the celestial sphere.

In our implementation of the Hough transform, we treat sky positions separately from frequencies and spin downs. In particular, we do not obtain the Hough histogram over the entire parameter space in one go, but rather for a given sky patch, a search frequency f_0 and a spin-down $\dot{f}$ value. These are the so-called *Hough maps* (HMs). Repeating this for every set of frequency and spin-down parameters and the different sky patches we wish to search over, we obtain a number of HMs. The collection of all these HMs represent our final histogram in parameter space.

The HMs could be produced by using a ''brute force'' approach, i.e., using all the peaks in the time-frequency plane. But there is an alternative way of constructing them. Let us define a *partial Hough map* (PHM) as being a Hough histogram, in the space of sky locations, obtained by performing the Hough transform using the peaks from a *single* SFT and for a *single* value of the intrinsic signal frequency and no spin down. This PHM therefore consists of only zeros and ones, i.e., the collection of the annuli corresponding to all peaks present in a single peak-gram. Then each HM can be obtained by summing the appropriate PHMs produced from different SFTs. If we add PHMs constructed by using the same intrinsic frequency, then the resulting HM refers to the same intrinsic frequency and no spin down. But note that the effect of a spin down in the signal is the same as having a time varying intrinsic frequency. This suggests a strategy to reuse PHMs computed for different frequencies at different times in order to compute the HM for a nonzero spin-down case.

Given the set of PHMs, the HM for a given search frequency f_0 and a given spin down $\dot{f}$ is obtained as follows: using Eq. (10) calculate the trajectory $\hat{f}(t)$ in the time-frequency plane corresponding to f_0 and $\dot{f}$. If the mid time stamps of the SFTs are $\{t_i\}$ $(i = 1 \ldots N)$, calculate $\hat{f}(t_i)$ and find the frequency bin that it lies in; select the PHM corresponding to this frequency bin. Finally, add all the selected PHMs to obtain the Hough map. This procedure is shown in Fig. 2.

This approach saves computations because it recognizes that the same sky locations can refer to different values of

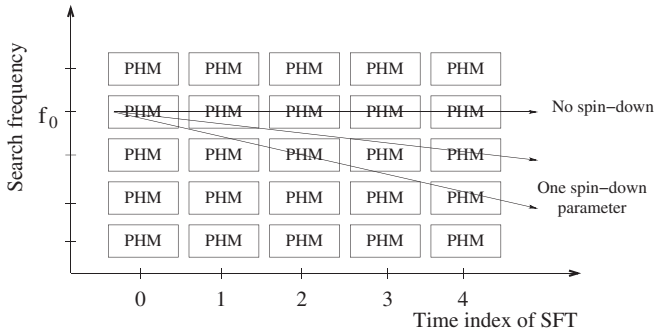

FIG. 2. A partial Hough map (PHM) is a histogram in the space of sky locations obtained by performing the Hough transform on a single SFT and for a given value of the instantaneous frequency. A total Hough map is obtained by summing over the appropriate PHMs. The PHMs to be summed over are determined by the choice of spin-down parameters which give a trajectory in the time-frequency plane.

frequency and spin down, and avoids having to redetermine such sky locations more than once. Another advantage of proceeding in this fashion is that we can use *look up tables* (LUTs) to construct the PHMs. The basic problem to construct the PHMs is that of drawing the annuli on the celestial sphere, or on the corresponding projected plane. The algorithm we use based on LUTs has proved to be more efficient than other methods we have studied, and this strategy is also employed by other groups [51].

A LUT is an array containing the list of borders of all the possible annuli, for a given value of $\mathbf{v}$ and $\hat{f}$, clipped on the sky-patch we use. Therefore it contains the coordinates of the points belonging to the borders that intersect the sky-patch, in accordance to the tiling we use, together with information to take care of edge effects. As described in [23], it turns out that the annuli are relatively insensitive to the value of the search frequency and, once a LUT has been constructed for a particular frequency, it can be reused for a large number of neighboring frequencies thus allowing for computational savings. The value of $\mathbf{v}$ used to construct the LUTs corresponds to the average velocity of the detector in the SSB frame during the 30 minutes interval of the corresponding SFT.

In fact, the code is further sped up by using *partial Hough map derivatives* (PHMDs) instead of the PHMs, in which only the borders of the annuli are indicated. A PHMD consists of only ones, zeros, and minus ones, in such a way that by integrating appropriately over the different sky locations one recovers the corresponding PHM. This integration is performed at a later stage, and just once, after summing the appropriate PHMDs, to obtain the final Hough map.

In the pipeline, we loop over frequency and spin-down values, taking care to update the set of PHMDs currently used, and checking the validity of the LUTs. As soon as the LUTs are no longer valid, the code recomputes them again together with the sky grid. Statistical analyses are per-

formed on the Hough maps in order to compress the output size. These include finding the maximum, minimum, mean and standard deviation of the number counts for each individual map, the parameters of the loudest event, and also of all the candidates above a certain threshold. We also record the maximum number count per frequency analyzed, maximized over all spin-down values and sky locations, and a histogram of the number counts for each 1 Hz band. The schematic search pipeline is shown in Fig. 3.

As a technical implementation detail, the search is performed by dividing the 200 Hz frequency band into smaller bands of 1 Hz and distributed using Condor [52] on a computer cluster. Each CPU analyzes a different 1 Hz band using the same pipeline (as described in Fig. 3). The code itself takes care to read in the proper frequency band from the SFTs. This includes the search band plus an extra interval to accommodate for the maximum Doppler shift, spin down, and the block sized used by the running median. The analysis described here was carried out on the Merlin cluster at AEI [53]. The full-sky search for the entire S2 data from the three detectors distributed on 200 CPUs on Merlin lasted less than half a day.

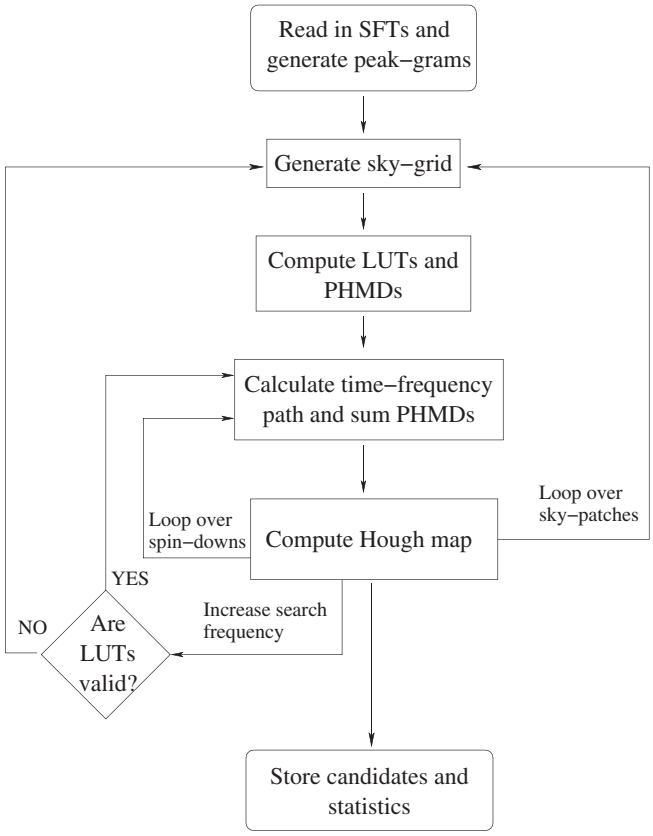

FIG. 3. The schematic of the analysis pipeline. The input data are the SFTs and the search parameters. The first step is to select frequency bins from the SFTs and generate the peak-grams. Then, the Hough transform is computed for the different sky patches, frequencies and spin-down values, thus producing the different Hough maps. The search uses LUTs that are computed for a given tiling of the sky-patch. The sky grid is frequency dependent, but it is fixed for the frequency range in which the LUTs are valid. Then, a collection of PHMDs is built, and for each search frequency f_0 and given spin down $\dot{f}$, the trajectory in the time-frequency plane is computed and the Hough map obtained by summing and integrating the corresponding PHMDs. The code loops over frequency and spin-down parameters, updating the sky grid and LUTs whenever required. Statistical analyses are performed on each map in order to reduce the output size.

The software used in the analysis is available in the LIGO Scientific Collaboration's CVS archives (see [54]), together with a suite of test programs, especially for visualizing the Hough LUTs. The full search pipeline has also been validated by comparing the results with independently written code that implements a less efficient but conceptually simpler approach, i.e., for each point in parameter space $(f_0, \dot{f}, \mathbf{n})$, it finds the corresponding pattern in the time-frequency plane and sums the corresponding selected frequency bins.

D. Number counts from L1, H1 and H2

In the absence of a signal, the distribution of the Hough number count ideally is a binomial distribution. Environmental and instrumental noise sources can excite the optically sensed cavity length, or get into the output signal in some other way, and show up as spectral disturbances, such as lines. If no data conditioning is applied, line interference can produce an excess of number counts in the Hough maps and mask signals from a wide area in the sky. Figure 4 shows the comparison of the theoretical binomial distribution Eq. (12) with the distributions that we obtain experimentally in two bands: 206–207 Hz and 343–344 Hz. The first band contains very little spectral disturbances while the second band contains some violin modes. As shown in Fig. 4, the Hough number count follows the expected binomial distribution for the *clean* band while it diverges from the expected distribution in the presence of strong spectral disturbances, such as the violin modes in this case. We have verified good agreement in several different frequency bands that were free of strong spectral disturbances.

The sources of the disturbances present in the S2 data are mostly understood. They consist of calibration lines, broad 60 Hz power line harmonics, multiples of 16 Hz due to the data acquisition system, and a number of mechanical resonances, as, for example, the violin modes of the mirror suspensions [36,55]. The 60 Hz power lines are rather broad, with a width of about ± 0.5 Hz, while the calibration lines and the 16 Hz data acquisition lines are confined to a single frequency bin. A frequency comb is also present in the data, having fundamental frequency at 36.867 Hz for L1, 36.944 Hz for H1 and 36.975 Hz for H2, some of them accompanied with side lobes at about 0.7 Hz, created by up-conversion of the pendulum modes of some core optics,

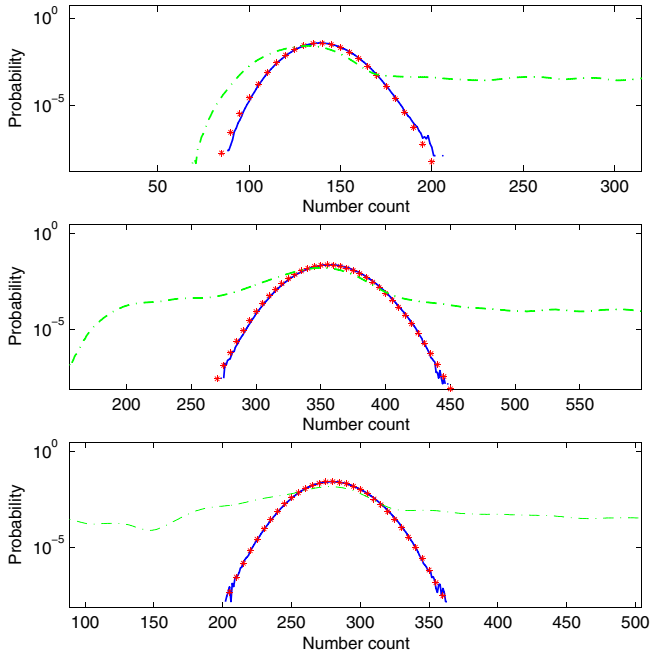

FIG. 4 (color online). Top: Graph of the L1 number-count discrete probability distribution: the solid curve corresponds to the number-count distribution obtained for the band between 206–207 Hz, the dash-dot curve to the number-count distribution obtained for the 343–344 Hz band, that contains violin modes, and in asterisks the theoretical expected binomial distribution for 687 SFTs and a peak selection probability of 20%. Middle: the H1 number-count distribution for 1761 SFTs. Bottom: the H2 number-count distribution for 1384 SFTs.

but these were only present (or at least prominent) in H1 and H2. The sources of these combs were synthesized oscillators used for phase modulation that were later replaced by crystal oscillators. In addition to the above disturbances, we also observe a large number of multiples of 0.25 Hz. While this comb of lines is strongly suspected to be instrumental, its physical origin is unknown. In Table I we summarize the list of known spectral disturbances in the three interferometers during the S2 run.

After analyzing all the data we discard the number counts from all those frequencies that could be contaminated by a *known* instrumental artifact. Thus, we exclude every frequency bin which is affected by the spectral disturbances including the maximum possible Doppler broadening of these lines; thus, for a known spectral disturbance at a frequency f and width Δf, we exclude a frequency range $\pm(vf/c + \Delta f/2)$ around the frequency f. We also exclude from our analysis the frequency band 342–348 Hz for L1 and H1 and 342–351 Hz for H2 because they contain many violin modes. The net effect of this vetoing strategy is that we consider only 67.1% for L1, 66.8% for H1 and 65.6% for H2 of the full 200 Hz range.

Figures 5–7 show the maximum Hough number count n_k^* obtained in each of the 360 000 different frequency bins f_k analyzed, maximized over all sky locations and spin-

TABLE I. List of known spectral disturbances in the three interferometers during the S2 run used as a frequency veto in the 200–400 Hz band. f refers to the central frequency and Δf to the full width of the lines. Lines denoted with $\Delta f = 0.0$ Hz are those in which the line width is much smaller than the associated maximum Doppler broadening of the line. This ranges from ~0.04 Hz at a frequency of 200 Hz up to ~0.08 Hz at 400 Hz.

L1		H1		H2	
f (Hz)	Δf (Hz)	f (Hz)	Δf (Hz)	f (Hz)	Δf (Hz)
0.250[a]	0.00	0.250[a]	0.00	0.250[a]	0.00
16.000[a]	0.00	16.000[a]	0.00	16.000[a]	0.00
60.000[a]	1.00	60.000[a]	1.00	60.000[a]	1.00
221.200	0.01	221.665	0.01	221.850	0.02
258.080	0.04	257.875	0.02	258.830	0.01
294.935	0.01	258.610	0.02	295.070	0.02
331.810	0.01	259.340	0.02	295.670	0.02
345.000	6.00	294.820	0.01	295.800	0.04
368.670	0.02	295.560	0.00	295.930	0.02
		296.300	0.00	296.530	0.02
		331.790	0.00	323.300	0.00
		332.490	0.00	323.870	0.04
		333.200	0.00	324.000	0.04
		335.780	0.14	324.130	0.04
		336.062	0.00	324.700	0.00
		339.000	0.02	332.800	0.00
		339.720	0.02	335.120	0.02
		345.000	6.00	335.590	0.02
		365.500	0.02	341.615	0.01
		368.690	0.00	346.500	9.00
		369.430	0.00	349.202	0.00
		370.170	0.01		

[a]Indicates all higher harmonics are present.

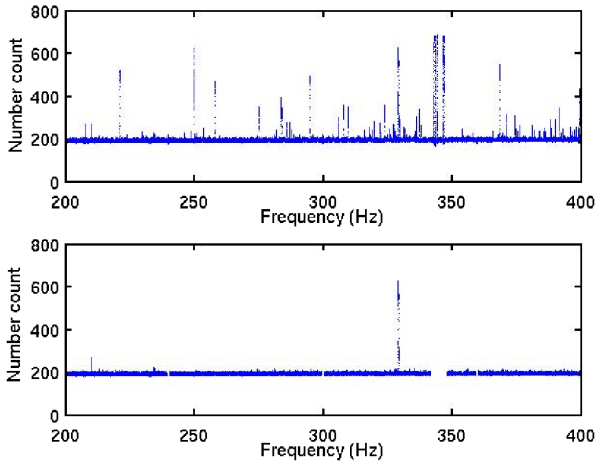

FIG. 5 (color online). Graph of the L1 maximum number count $n_k^\star$ for every analyzed frequency f_k, maximized over all spin-down values and sky locations. The top figure corresponds to the raw output from the Hough transform in which many outliers are clearly visible. The bottom figure corresponds to the same data after vetoing the frequency bands contaminated by known instrumental noise. See Appendix B for details on outliers.

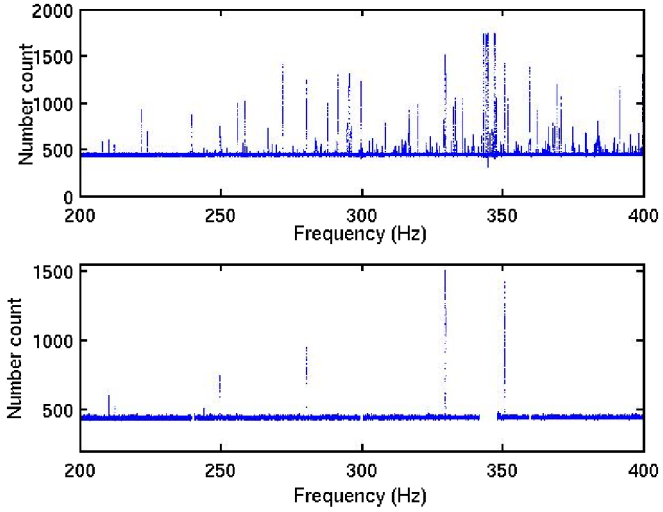

FIG. 6 (color online). Graph of the H1 maximum number count n_k^* versus frequency f_k as in Fig. 5.

down values for the three detectors. In those figures we compare the results of the search before and after applying the frequency veto, showing how the spectral disturbances impact on the number counts. This can be seen more clearly in Fig. 8, where we plot the histograms of these maximum Hough number counts n_k^* before and after applying the frequency veto.

These values of $n_k^\star$ obtained after removing the known outliers can be easily seen to be consistent with what we would expect for Gaussian stationary noise. Assuming that the maximum occurs only once, the expected value of $n_k^\star$ should be consistent with a false alarm rate of $1/m_{\rm bin}(f_k)$, where $m_{\rm bin}(f_k)$ is the total number of templates at a frequency f_k. $m_{\rm bin}(f_k)$ is frequency dependent, ranging from $\sim 1.1 \times 10^6$ at 200 Hz to $\sim 4.2 \times 10^6$ at 400 Hz. Thus,

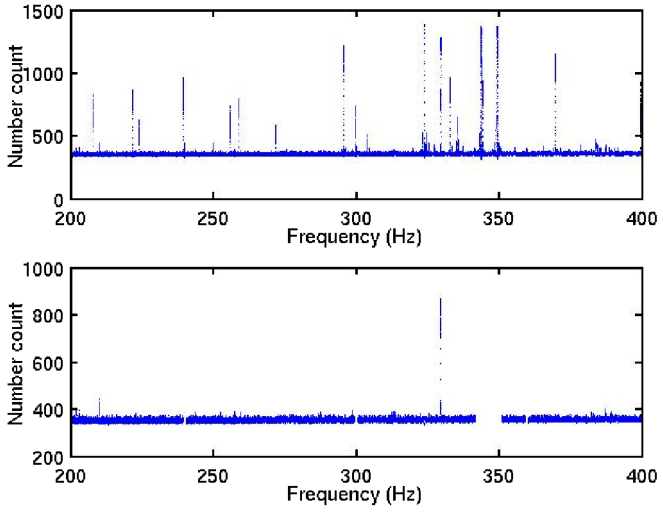

FIG. 7 (color online). Graph of the H2 maximum number count n_k^* versus frequency f_k as in Fig. 5.

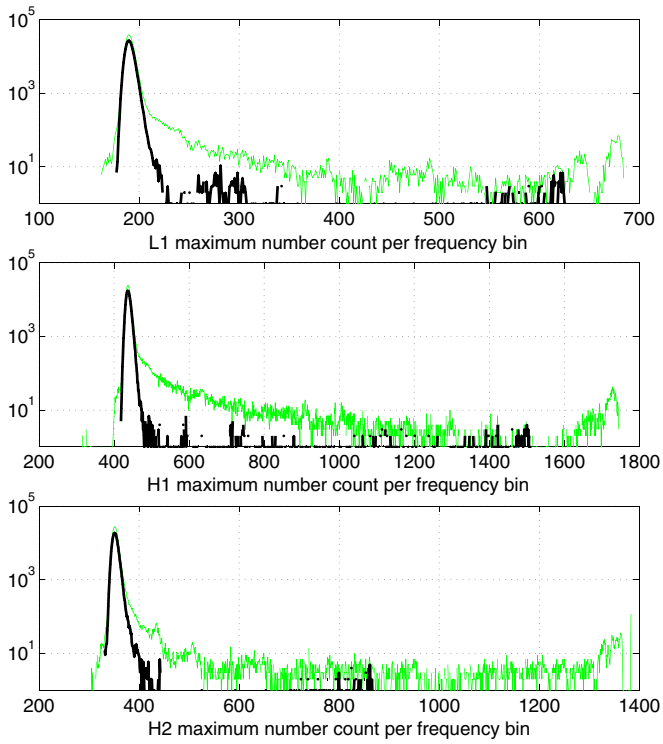

FIG. 8 (color online). Histograms of the maximum Hough number count n_k^* for the three detectors. The light curve corresponds to the raw output and the dark thick curve corresponds to the same output after vetoing the contaminated frequencies.

$n_k^\star$ should be consistent with a false alarm rate of $\alpha_{400} = 2.4 \times 10^{-7}$ at 400 Hz, up to $\alpha_{200} = 9.5 \times 10^{-7}$ at 200 Hz. In the case of Gaussian stationary noise, the expectation value of n_k^* should therefore be similar to $n_{\rm th}$ [defined by Eq. (16)] for a false alarm $\alpha = 1/m_{\rm bin}(f_k)$. Table II compares the mean $\langle n_k^\star \rangle$ of the maximum Hough number count before and after vetoing, with $n_{\rm th}(\alpha)$ at different false alarm rates. After vetoing, we observe that $\langle n_k^* \rangle$ lies within the interval $[n_{\rm th}(\alpha_{200}), n_{\rm th}(\alpha_{400})]$, and the standard deviation std(n_k^*) is also greatly reduced. This indicates the consistency of the observed values of $n_k^\star$ with ideal noise.

As can be seen in Figs. 5–7, a few outliers remain after applying the frequency veto described above. The fact that these outliers have such small false alarm probabilities makes it very unlikely they were drawn from a parent

TABLE II. Comparison between the statistics of the maximum Hough number count n_k^* (before and after the frequency veto) and $n_{\rm th}(\alpha)$ at different false alarm rates $\alpha_{200} = 9.5 \times 10^{-7}$ and $\alpha_{400} = 2.4 \times 10^{-7}$ for the three detectors.

			Before veto		After veto	
	$n_{\rm th}(\alpha_{200})$	$n_{\rm th}(\alpha_{400})$	$\langle n_k^* \rangle$	std(n_k^*)	$\langle n_k^* \rangle$	std(n_k^*)
L1	188.8	191.7	194.8	32.3	191.4	10.2
H1	435.8	440.3	452.0	94.0	439.8	29.5
H2	350.6	354.6	360.6	72.3	353.2	11.3

Gaussian distribution. However they could also be due to spectral disturbances (line noise) that mimic the time-frequency evolution of a pulsar for a certain location in the sky. If these outliers are due to gravitational signals they should show up in the different detectors. By performing a simple coincidence analysis in frequency, the only candidate that remains is the one at 210.360 Hz. But this has been ruled out since it seems to be associated to multiples of 70.120 Hz produced by a VME (VERSA module Eurocard) controller hardware used during S2. Since the other outliers are not coincident among the three detectors, there is no evidence for a detection. We refer the interested reader to Appendix B for further details.

As explained in Sec. I, the ultimate goal for wide parameter-space searches for continuous signals over large data sets is to employ hierarchical schemes which alternate coherent and semicoherent techniques. The Hough search would then be used to select candidates in the parameter space to be followed up. The way in which those candidates would be selected is the following: fix the number of candidates to follow up, determine the false alarm rate, and set the corresponding threshold on the Hough number count. Not all the candidates selected in this manner will correspond to real gravitational-wave signals, but they will point to interesting areas in parameter space.

The analysis presented here is a very important step forward in this direction. However, given the limited sensitivity its relevance mostly rests in the demonstration of this analysis technique on real data. In what follows we will thus concentrate on setting upper limits on the amplitude h_0 of continuous gravitational waves emitted at different frequencies.

VII. UPPER LIMITS

We use a frequentist method to set upper limits on the amplitude h_0 of the gravitational-wave signal. Our upper limits refer to a hypothetical population of isolated spinning neutron stars which are uniformly distributed in the sky and have a spin-down rate $\dot{f}$ uniformly distributed in the range $(-\dot{f}_{\max}, 0)$. We also assume uniform distributions for the parameters $\cos\iota \in [-1, 1]$, $\psi \in [0, 2\pi]$, and $\Phi_0 \in [0, 2\pi]$. As before, the frequency range considered is 200–400 Hz.

The upper limits on h_0 emitted at different frequencies are based on the highest number count, the *loudest event*, registered over the entire sky and spin-down range at that frequency. Furthermore, we choose to set upper limits not on each single frequency but on a set of frequency values lying within the same 1 Hz band and thus are based on the highest number count in each frequency band. In every 1 Hz band the loudest event is selected excluding all the vetoed frequencies of Table I.

Let $n^\star$ be the loudest number count measured from the data. The upper limit h_0^C on the gravitational-wave amplitude, corresponding to a confidence level C, is the value such that had there existed in the data a real signal with an amplitude greater than or equal to h_0^C, then in an ensemble of identical experiments with different realizations of the noise, some fraction C of trials would yield a maximum number count equal to or bigger than $n^\star$. The upper limit h_0^C corresponding to a confidence level C is thus defined as the solution to this equation:

$$\text{Prob}\,(n \geq n^\star | h_0^C) = \sum_{n=n^\star}^{N} p(n|h_0^C) = C, \qquad (23)$$

where $p(n|h_0)$ is the number-count distribution in the presence of a signal with amplitude h_0 and averaged over all the other parameters; note that $p(n|h_0)$ is different from the distribution $p(n|h)$ discussed before Eq. (12) which was relevant for a targeted search when *all* signal parameters are known. We choose to set upper limits at a confidence level of $C = 95\%$; $h_0^{95\%}$ denotes the 95% confidence upper limit value.

Given the value of $n^\star$, if the distribution $p(n|h_0)$ were known, it would be a simple matter to solve Eq. (23) for h_0^C. In the absence of any signal, this distribution is indeed just a binomial, and as exemplified in Fig. 4, this is also what is observed in practice. If a signal were present, the distribution may not be sufficiently close to binomial because the quantity λ defined in Eq. (14) varies across the SFTs due to nonstationarity in the noise and the amplitude modulation of the signal for different sky locations and pulsar orientations. In addition, now we must also consider the random mismatch between the signal and template (in relation to the parameter-space resolution used in the analysis) which causes an additional reduction in the effective signal-to-noise ratio for the template. For these reasons, we measure $p(n|h_0)$ by means of a series of software injections of fake signals in the real data. Figure 9 shows four distributions for different h_0 values obtained by Monte Carlo simulations. While for low signal amplitudes the distribution is close to the ideal binomial one, the distribution diverges from it at higher amplitudes, thereby illustrating the complexity of the number-count distribution for sufficient large h_0.

Our strategy for calculating the 95% upper limits is to find $p(n|h_0)$ for a wide range of h_0 values, then to get the corresponding confidence levels $C(h_0)$, and find the two values of h_0 which enclose the 95% confidence level. The 95% upper limit is approximated by a linear interpolation between these values. We then refine this reduced range of h_0 values by further Monte Carlo simulations until the desired accuracy is reached. This is described in Fig. 10.

The parameters of the fake injected signals are drawn from the population described above, and we ensure that the frequency does not lie within the excluded bands. The data with the injections are searched using the search pipeline used for the actual analysis. For computational efficiency, for each injected signal we find the number count using only the 16 nearest templates, and choose the

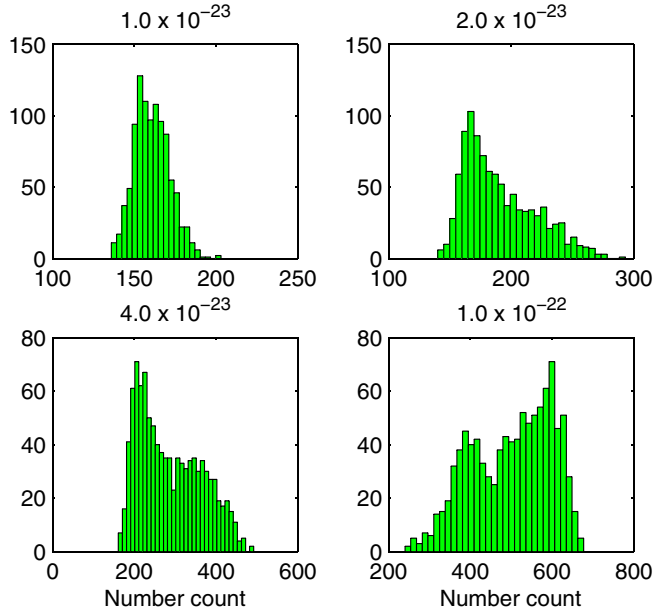

FIG. 9 (color online). Histograms of the Hough number-count distribution (in arbitrary units) for L1 using 1000 injected signals randomly distributed over the whole sky within the band 200–201 Hz for different h_0 values. The largest number count for the search in that band was 202. The confidence level associated with the different h_0 values are: 0.1% for 1.0×10^{-23}, 30.5% for 2.0×10^{-23}, 87.0% for 4.0×10^{-23}, and 1 for 1.0×10^{-22}.

one yielding the largest number count. The SFT data in the different frequency bins in a 1 Hz band (of order $N \times 1800$ bins that get combined differently for the different time-frequency patterns) can be considered as different realizations of the same random process. Therefore it is reasonable to assume that the normalized histogram of the largest number count represents the discrete probability distribution $p(n|h_0)$.

The most computationally intensive part of this Monte Carlo scheme is the generation of the artificial signals. The computational costs can be greatly reduced by estimating $p(n|h_0)$ for different h_0 values in one go: for each individual artificial signal, we generate a set of SFTs containing only one noiseless signal of a given amplitude. These SFTs can be scaled by an appropriate numerical constant to obtain a set of SFTs containing signals with different amplitudes, which are then added to the noise SFTs. The disadvantage of doing this is that the different signals obtained by rescaling the amplitude this way are not statistically independent since all the other signal parameters are identical. We must ensure that we have a sufficiently large number of trials so that the error in the final upper limit is sufficiently small.

We have found empirically that 5000 injections per band are sufficient to get upper limits accurate to within 3%; see Figs. 11–13. Figure 11 shows the confidence level as a function of the signal amplitude h_0 for L1 data within the

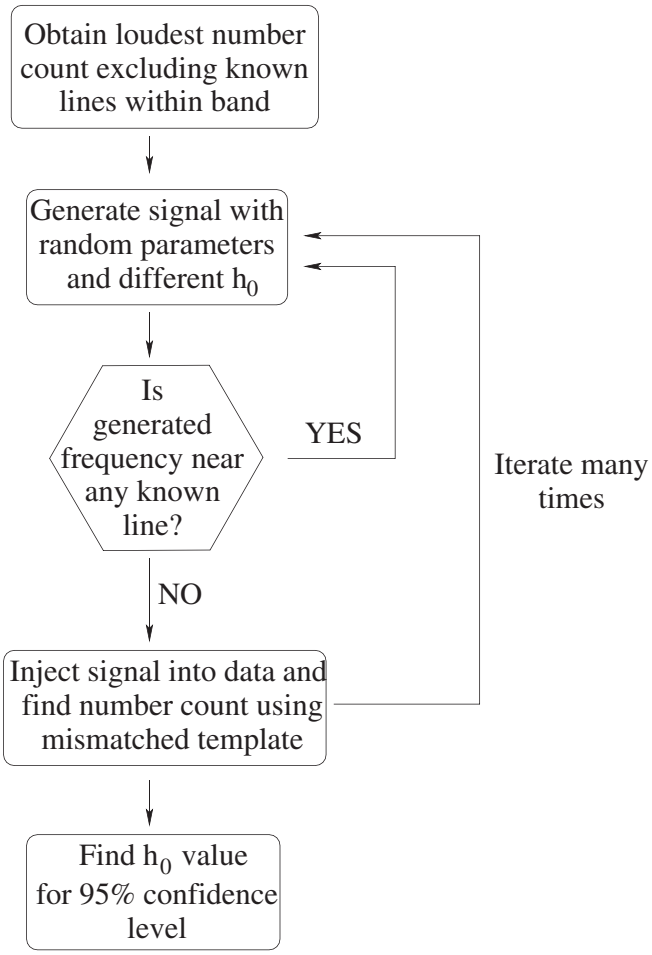

FIG. 10. Pipeline description of the Monte Carlo simulation to determine the upper limit values $h_0^{95\%}$. We inject randomly generated fake pulsar signals with fixed amplitude (the other parameters are drawn from a suitable uniform distribution) into the real data and measure the value of h_0 required to reach a 95% confidence level.

band 200–201 Hz. The solid line corresponds to our most accurate simulations using 100 000 injections. The two dashed lines correspond to two different simulations both using 5000 injected signals and the dotted line corresponds to simulation with 3000 injections. In each case, confidence levels for different h_0 values are calculated by simply rescaling the signal as described above. This means that all the points in each individual curve in Fig. 11 are statistically biased in a similar way, and this explains why the curves in this figure do not intersect. To estimate the error in the 95% upper limit, we generate several of these curves for a fixed number of injections and we measure the error in $h_0^{95\%}$ with respect to the accurate reference value obtained using 100 000 injections. For this particular band and detector, we find that for 5000 injections, the error in the upper limit is at most 0.1×10^{-23}, corresponding to a relative error of 2.2%. This experiment has been repeated for several 1 Hz bands and many simulations, and in all of

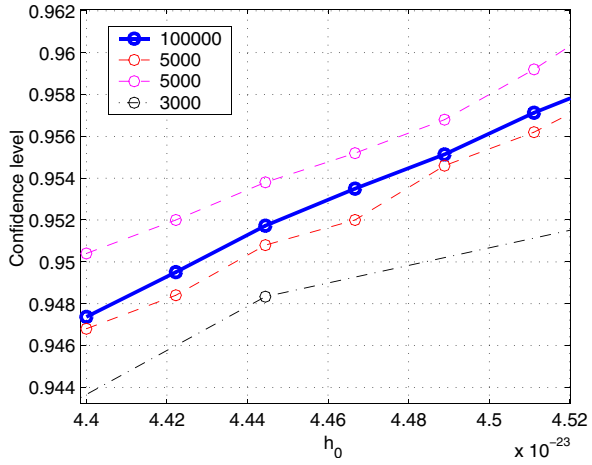

FIG. 11 (color online). Confidence level as a function of the signal amplitude h_0 for different Monte Carlo simulations for L1 within the band 200–201 Hz. The solid line corresponds to a simulation using 100 000 injected signals. The two dashed lines correspond to 5000 injected signals. The dash-dotted line corresponds to 3000 signal injections. The $h_0^{95\%}$ upper limit for this band using and these two simulations with 5000 injections has a maximum absolute error of the order of 0.02×10^{-23} corresponding to a relative error smaller than 0.5%. In the case of using only 3000 injections the error increases to the 1.5% level.

them 5000 injections were enough to ensure an accuracy of better than 3% in the $h_0^{95\%}$. We also have found that with 5000 injections, using amplitudes equal to the most accurate upper limit $h_0^{95\%}$ obtained from 100 000 injections, the confidence level are within (94.5–95.5)%. See Figs. 12 and 13.

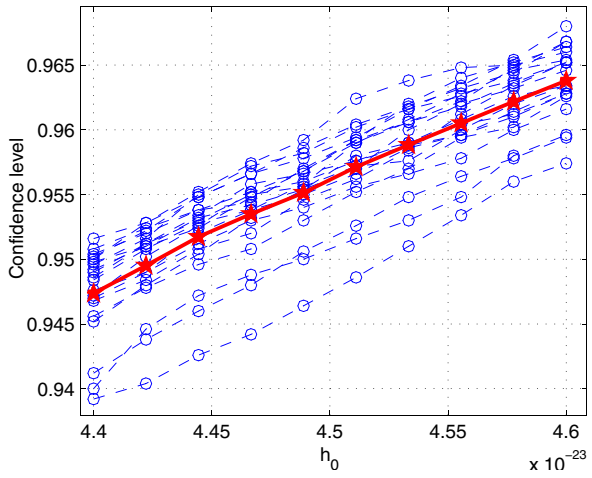

FIG. 12 (color online). Confidence level as a function of the signal amplitude h_0 for different Monte Carlo simulations for L1 within the band 200–201 Hz. The solid thick line corresponds to a simulation using 100 000 injected signals. The other 20 lines correspond to simulations with 5000 injected signals. The $h_0^{95\%}$ upper limit for this band using 5000 injections has a maximum absolute error of the order of 0.1×10^{-23} corresponding to a relative error of 2.2%.

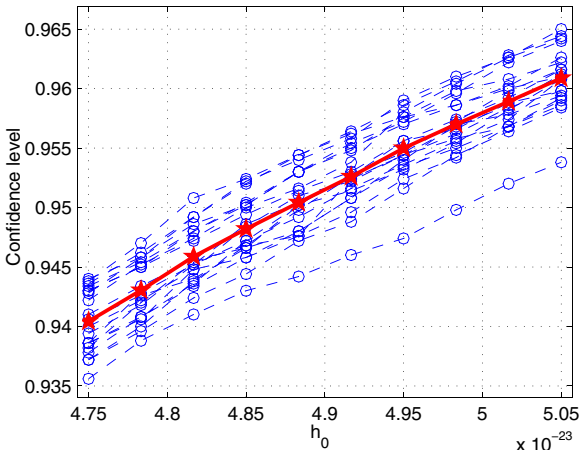

FIG. 13 (color online). Confidence level as a function of the signal amplitude h_0 for different Monte Carlo simulations for H1 within the band 259–260 Hz. The solid thick line corresponds to a simulation using 100 000 injected signals. The other 20 lines correspond to simulations with 5000 injected signals. The $h_0^{95\%}$ upper limit for this band using 5000 injections has a maximum absolute error of the order of 0.1×10^{-23} corresponding to a relative error of 2.0%.

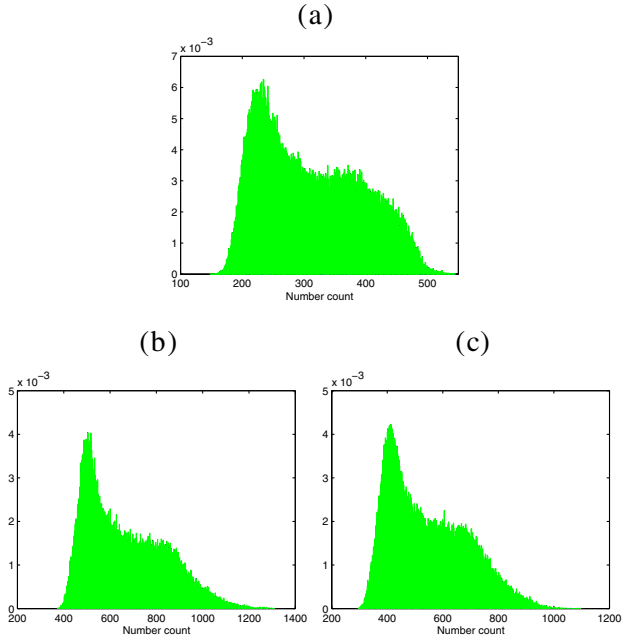

FIG. 14 (color online). Measured Hough number-count probability distribution for $p(n|h_0^{\text{inject}})$ from the Monte Carlo simulations using 100 000 injected signals. (a) L1 within the band 200–201 Hz, with a h_0^{inject} of 4.422×10^{-23} corresponding to a confidence level of 94.95%. The $n^\star$ value for this band was 202. (b) H1 within the band 259–260 Hz, with a h_0^{inject} of 4.883×10^{-23} corresponding to a confidence level of 95.04%. The $n^\star$ value for this band was 455. (c) H2 within the band 258–259 Hz, with a h_0^{inject} of 8.328×10^{-23} corresponding to a confidence level of 95.02%. The $n^\star$ value for this band was 367.

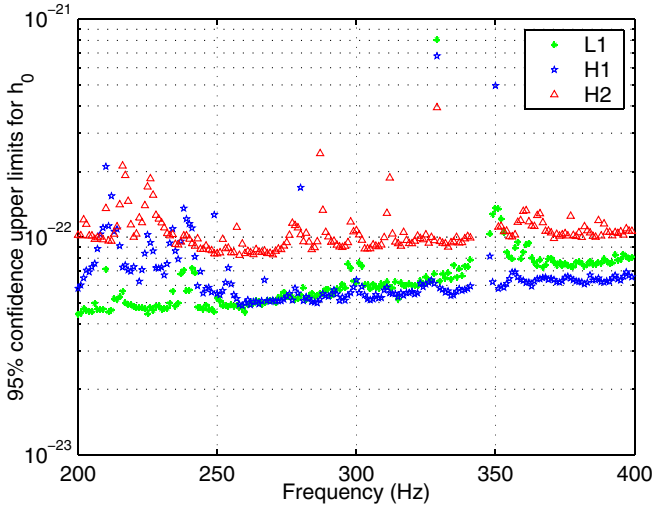

FIG. 15 (color online). The 95% confidence upper limits on h_0 over the whole sky and different spin-down values in 1 Hz bands.

Figure 14 shows the probability distribution for $p(n|h_0^{\text{inject}})$ for L1, H1 and H2 in a 1 Hz band measured from 100 000 randomly injected signals, with a signal amplitude h_0^{inject} *close* to the 95% confidence level; note that these h_0^{inject} values are not the 95% upper limits because they correspond to different confidence levels. This illustrates, yet again, that the true number-count distributions are far from the ideal binomial distribution. It is also interesting to notice that, among all the templates used to analyze each injection, the nearest template, in the normal Euclidean sense, corresponds to the template providing the highest number count only 18% of the time. This effect is due to the noise contribution and the fact that the matching-amplitude, described by the parameter-space metric is highly anisotropic [56].

The 95% confidence upper limits on h_0 for each 1 Hz band using all the data from the S2 run are shown in Fig. 15. As expected, the results are very similar for the L1 and H1 interferometers, but significantly worse for H2. The typical upper limits in this frequency range for L1 and H1 are mostly between ~ 4–9×10^{-23}, typically better at lower frequencies. The most stringent upper limit for L1 is 4.43×10^{-23} which occurs within 200–201 Hz, largely reflecting the lower noise floor around 200 Hz. For H1 it is 4.88×10^{-23}, which occurs in the frequency range 259–260 Hz, and for H2 it is 8.32×10^{-23}, which occurs in the frequency range 258–259 Hz. The values of the most stringent upper limits on $h_0^{95\%}$ have been obtained using 100 000 injections in the most sensitive 1 Hz bands. The upper limits are significantly worse in bands lying near the known spectral disturbances, especially near the violin modes.

In Table III we summarize the best upper limits on $h_0^{95\%}$ and we compare them with the theoretical values h_0^{exp} we would expect for a directed search using a perfectly matched template, as given by Eq. (17). Here we take a false dismissal rate of 5% and the false alarm rate associated to the loudest number count in that band. In those three bands the ratio $h_0^{95\%}/h_0^{\text{exp}}$ is about 1.8.

These $h_0^{95\%}$ results are also about a factor of 2.6 worse than those predicted by Fig. 1 which corresponds to a directed search using a perfectly matched template, and with a false alarm rate of 1% and a false dismissal rate of 10%. Of these 2.6, a factor ~ 1.5 is due to the use of different values of the false alarm and false dismissal rate, and a factor ~ 1.8 because the $p(n|h_0)$ distribution does not correspond to the ideal binomial one for values of h_0 distinct from zero.

From the upper limits on h_0, using Eq. (5), we can derive the distance covered by this search. This is shown in Fig. 16 assuming $\epsilon = 10^{-6}$ and $I_{zz} = 10^{45}\ \text{g cm}^2$. The maximum reach is 2.60 pc for H1 which occurs within 395–396 Hz. For L1 the maximum reach is 2.15 pc and 1.62 pc for H2. This value of $\epsilon = 10^{-6}$ corresponds to the maximum expected ellipticity for a regular neutron star, but ellipticities from more exotic alternatives to neutron stars may be larger [10], e.g., solid strange quark stars for which $\epsilon_{\text{max}} \approx 10^{-4}$. Therefore the astrophysical reach for these exotic stars could be better by a factor 100.

TABLE III. Best all-sky upper limits obtained on the strength of gravitational waves from isolated neutron stars. The $h_0^{95\%}$ values have been obtained using 100 000 injections in the best 1 Hz bands f_{band} for the three detectors. $\langle S_n(f_{\text{band}})\rangle$ is the average value of noise in that 1 Hz band excluding the vetoed frequencies, N the number of SFTs available for the entire S2 run, n^* is the loudest number-count measured from the data in that band, α^* is the corresponding false alarm assuming Gaussian stationary noise derived from Eq. (16), m_{1Hz} the number of templates analyzed in that 1Hz band, the quantify S^{exp} is defined by Eq. (18) using the values $\alpha = \alpha^*$ and $\beta = 0.05$. h_0^{exp} is the theoretical expected upper limit from searches with perfectly matched templates assuming the ideal binomial distribution for $p(n|h_0)$ defined by Eq. (17), therefore ignoring also the effects of the different sensitivity at different sky locations and pulsar orientations.

Detector	Best $h_0^{95\%}$	f_{band} (Hz)	$\langle S_n(f_{\text{band}})\rangle$ (Hz^{-1})	N	n^*	α^*	$m_{1\ \text{Hz}}$	S^{exp}	h_0^{exp}
L1	4.43×10^{-23}	200–201	1.77×10^{-43}	687	202	8.94×10^{-10}	1.88×10^{9}	5.4171	2.41×10^{-23}
H1	4.88×10^{-23}	259–260	3.53×10^{-43}	1761	455	1.77×10^{-9}	3.11×10^{9}	5.3381	2.67×10^{-23}
H2	8.32×10^{-23}	258–259	1.01×10^{-42}	1384	367	2.25×10^{-9}	3.11×10^{9}	5.3098	4.78×10^{-23}

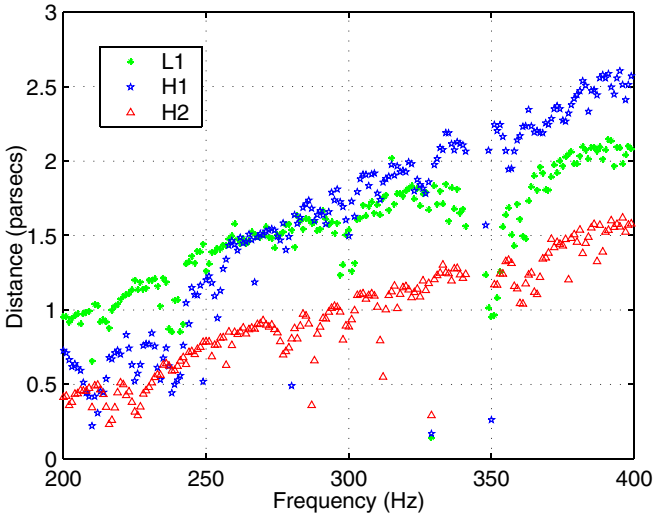

FIG. 16 (color online). Astrophysical reach covered by the search. The curves show the estimated distance d out to which signals from isolated gravitational-wave pulsars could be detected in our S2 data set derived from the upper limits on h_0. This plot assumes $\epsilon = 10^{-6}$ and $I_{zz} = 10^{45}$ g cm^2.

VIII. HARDWARE INJECTIONS

Two artificial pulsar signals, based on the waveforms given in Eqs. (2)–(4) and (6) were injected into all three LIGO interferometers for a period of 12 h towards the end of the S2 run. These injections were designed to give an end-to-end validation of the search pipeline. The waveforms were added digitally into the interferometer length sensing and control system (responsible for maintaining a given interferometer on resonance), resulting in a differential length dither in the optical cavities of the detector. We denote the two pulsars P1 and P2; their parameters are given in Table IV.

TABLE IV. Parameters of the two hardware injected pulsars. See Eqs. (2)–(4) and (6) for the definition of the parameters. RA and Dec are the right ascension and declination in equatorial coordinates. T_0 is the GPS time in the SSB frame in which the signal parameters are defined. h_0 is the amplitude of the signal according to the strain calibration used at the time of the injections.

	P1	P2
f_0 (Hz)	1279.123 456 789 012	1288.901 234 567 890 123
$\dot{f}$ (Hz-s^{-1})	0	-10^{-8}
RA (rad)	5.147 162 131 9	2.345 678 901 234 567 890
Dec (rad)	0.376 696 024 6	1.234 567 890 123 45
ψ (rad)	0	0
$\cos\iota$	0	0
Φ_0 (rad)	0	0
T_0 (sec)	733 967 667.026 112 310	733 967 751.522 490 380
h_0	2×10^{-21}	2×10^{-21}

The data corresponding to the injection period have been analyzed using the Hough transform and the same search code as described in Sec. VI C. As before, the input data consists of 30 min long SFTs. The number of SFTs available are 14 for L1, 17 for H1 and 13 for H2. As in Sec. VI B, the frequency resolution is $1/1800$ s, the sky resolution is given by Eq. (22), and $\delta\dot{f} = 1/(T_{\rm obs}T_{\rm coh})$. Since the total effective observation time is somewhat different for the three detectors, we get $\delta\dot{f} = -2.286\,24 \times 10^{-8}$ Hz-s^{-1} for L1, $-1.770\,24 \times 10^{-8}$ Hz-s^{-1} for H1, and $-1.935\,33 \times 10^{-8}$ Hz-s^{-1} for H2. As before, for each intrinsic frequency we analyze 11 different spin-down values. The portion of sky analyzed was of 0.5 radians $\times$0.5 radians around the location of the two injected signals.

Figure 17 shows the Hough maps corresponding to the nearest frequency and spin-down values to the injected ones. Although the presence of the signal is clearly visible, it is apparent that 12 h of integration time is not enough to

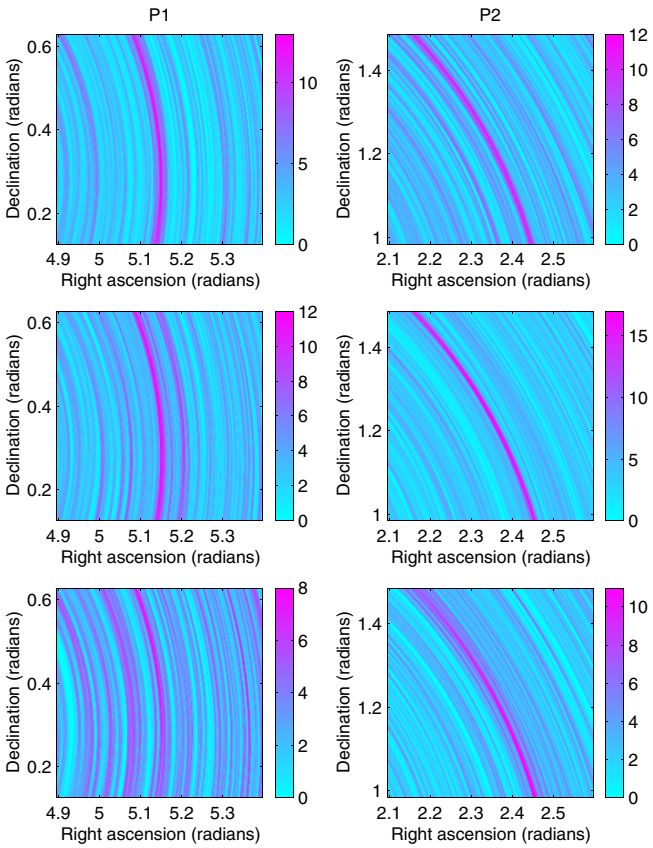

FIG. 17 (color online). Hough maps for the hardware injected signals. The left maps correspond to P1 and the right maps to P2. From top to bottom, the left maps correspond to L1, H1 and H2, for a frequency of 1279.123 333 Hz and zero spin down; the right maps to a frequency of 1288.901 11 Hz and zero, -1.77024×10^{-8} Hz-s^{-1}, and -1.93533×10^{-8} Hz-s^{-1} spin-down values, respectively. The location of the injected signals are close to the centers of these subplots.

identify the location of the source using the Hough transform. In particular if one looks at the Hough maps at mismatched frequencies and spin-down values, one can still identify annuli with very high number-counts, but appearing with a mismatched sky location.

For P2, the Hough maps corresponding to the closest values of frequency and spin down contain the maximum number count at the correct sky location. These maximum number counts are 12 for L1, 17 for H1, and 11 for H2. Notice that for L1 and H2 these numbers are smaller than the number of SFTs used.

For P1, the closest template to the signal parameters has a number count of 13 for L1, 12 for H1 and 8 for H2. The maximum number counts obtained in the search were 13 for L1, but 14 for H1, and 10 for H2. Those higher number counts occurred for several templates with a larger mismatch, for example, at 1279.123 333 Hz and $-5.310\,73 \times 10^{-8}$ Hz-s^{-1} for H1, and 1279.132 222 Hz and $-1.354\,73 \times 10^{-7}$ Hz-s^{-1} for H2. This is not surprising because we only compute the Hough maps at the Fourier frequencies $n \times 1/T_{\mathrm{coh}}$. In any case, both pulsar signals are unambiguously detected because these number counts are much bigger then the expected average number counts for pure noise.

IX. CONCLUSIONS

In this paper we use the Hough transform to search for periodic gravitational-wave signals. This is a semicoherent suboptimal method. Its virtues are computational efficiency and robustness. The search pipeline was validated using a series of unit tests and comparisons with independently written code. We then applied this method to analyze data from the second science run of all three LIGO interferometers. We also validated the search pipeline by analyzing data from times when two artificial pulsar signals were physically injected into the interferometer hardware itself. We show in this paper that the injected signals were clearly detected.

Our final results are all-sky upper limits on the gravitational-wave amplitude for a set of frequency bands. The best upper limits that we obtained for the three interferometers are given in Table V. The overall best upper limit is 4.43×10^{-23}. We searched the 200–400 Hz band and the spin-down space $\dot{f}_{\mathrm{max}} \leq 1.1 \times 10^{-9}$Hz-s^{-1}, and no vetoes were applied except for the list of ignored frequency bands that contain instrument line artifacts.

TABLE V. Best all-sky upper limits obtained on the strength of gravitational waves from isolated neutron stars. The $h_0^{95\%}$ values have been obtained using 100 000 injections in the best 1 Hz bands.

Detector	f_{band} (Hz)	$h_0^{95\%}$
L1	200–201	4.43×10^{-23}
H1	259–260	4.88×10^{-23}
H2	258–259	8.32×10^{-23}

Our best upper limit is 26 times worse than the best upper limit obtained for a *targeted* coherent search using the same data. This was an upper limit of 1.7×10^{-24} [16], achieved for PSR J1910-5959D. This is to be expected because we have performed not a targeted, but a wide parameter-space search. If we were to use the optimal $\mathcal{F}$-statistic method to perform a hypothetical search over the same parameter-space region as the Hough search in this paper, the number of templates required would be much larger for the same observation time: $\sim 10^{19}$ [20] instead of $\sim 10^{12}$. Thus, we would have to set a lower false alarm rate for this hypothetical search, and in the end, the sensitivity turns out to be roughly comparable to that of the Hough search. Note also that this hypothetical search is not computationally feasible for the foreseeable future.

From Eq. (1), in an optimistic scenario, we see that our best upper limits are approximately 1 order of magnitude larger than the strongest expected signals based on the statistics of the neutron star population of our galaxy. From Fig. 1, we see that with 1 yr of data at design sensitivity, the present method should gain us about one order or magnitude in sensitivity, thus enabling us to detect signals smaller than what is predicted by Eq. (1), at least for a certain frequency range. We can search for smaller signals either by increasing the number of coherent segments, N, or by increasing the coherent time baseline, T_{coh}. Since the number of segments is determined by the length of the data set, for a given amount of data one wants to make T_{coh} as large as possible. However, for the search pipeline presented in this paper, increasing T_{coh} was not possible due to the restriction on its value mentioned in Sec. V. This will be overcome by demodulating each short segment, taking into account the frequency and amplitude modulation of the signal. The optimal method will be to calculate the $\mathcal{F}$-statistic [57] for each segment. The time-frequency pattern will then no longer be given by (8) and (10) but by the master equation given in Ref. [23].

Since the wide parameter-space search for periodic gravitational-wave signals is computationally limited, there is also a limit on the maximum T_{coh} that can be used, given finite computing resources. Thus, a hierarchical strategy that combines fully coherent and semicoherent methods will be needed to achieve optimal results [21,22]. Our goal is to use the Hough transform as part of such a strategy. This is work in progress and the results will be presented elsewhere.

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 Educación y Ciencia, the John Simon Guggenheim Foundation, the David and Lucile Packard Foundation, the Research Corporation, and the Alfred P. Sloan Foundation. This document has been assigned LIGO Laboratory document number LIGO-P050013-00-R.

APPENDIX A: THE BIAS IN THE RUNNING MEDIAN

We are using a running median to estimate the noise floor in our SFTs. Thus, the value of S_n at a particular frequency bin can be estimated from the *median* of $|\tilde{x}_k|^2$ in w frequency bins around the frequency bin, where w is an integer and represents the window size of the running median. The reason for using a running median is to minimize the effect of large spectral disturbances, and to ensure that the presence of any pulsar signals will not bias our estimation of S_n. To carry this out in practice, we would like to know how the median can be used as an estimator of the mean. In this appendix, we answer this question assuming that the noise is Gaussian, so that the power is distributed exponentially.

Let x be a random variable with probability distribution $f(x)$. Let $F(x)$ denote the cumulative distribution function:

$$F(x) = \int_{-\infty}^{x} f(x')dx'. \tag{A1}$$

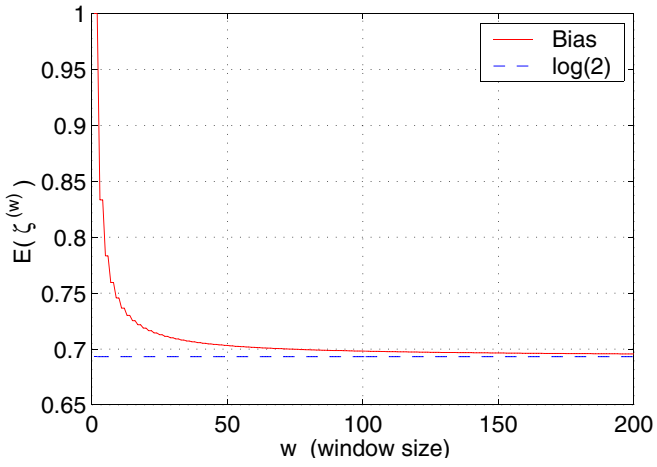

FIG. 18 (color online). Value of the running median bias $\zeta^{(w)}$ as a function of the window size w; $\zeta^{(w)}$ approaches log(2.0) for large w.

Let us draw w samples from this distribution and arrange them in increasing order: x_n $(n = 1 \ldots w)$. Define an integer k which is $(w-1)/2$ when w is odd, and $w/2$ when w is even. We define the median $\zeta^{(w)}$ as

$$\zeta^{(w)} = \begin{cases} x_{k+1} & \text{when } w \text{ is odd,} \\ \frac{1}{2}(x_k + x_{k+1}) & \text{when } w \text{ is even.} \end{cases} \tag{A2}$$

Consider first the case when w is odd. The distribution of $\zeta^{(w)}$ can be found as follows: $\zeta^{(w)}$ lies within the range $(x, x + dx)$ when k values are less than x, $w - k - 1$ values are greater than $x + dx$, and one value is in the range $(x, x + dx)$. The probability density for $\zeta^{(w)}$ is thus

$$g_w(x) = \binom{w}{k}(w-k)[F(x)]^k[1-F(x)]^{w-k-1}f(x). \tag{A3}$$

When w is very large, it can be shown that the distribution $g_w(x)$ approaches a Gaussian whose mean is equal to the

TABLE VI. List of outliers present in the Hough maps after applying the known instrumental frequency veto, for a false alarm rate of 10^{-13}. This corresponds to a threshold in the number count of 216 for L1, 480 for H1 and 390 for H2. For each outlier we quote the central frequency and the maximum number count. The triple coincidence at 210.36 Hz is a harmonic of the 70.12 Hz spectral disturbance described in the text.

Detector	Frequency (Hz)	Number count
L1	210.36	268
	234.50	224
	281.35	218
	329.34	626
	335.62	219
H1	210.36	596
	212.26	519
	244.14	507
	249.70	746
	280.48	949
	329.58	1510
	329.78	1227
	348.45	482
	350.60	1423
H2	202.18	402
	203.23	395
	210.36	443
	298.81	394
	329.69	867
	387.05	400
	389.40	391

population median $\hat{\zeta}$, and whose variance is proportional to $1/\sqrt{w}$ [58].

For Gaussian noise, the normalized power ρ follows the exponential distribution, i.e. $f(x) = e^{-x}$ for $x \geq 0$ and $f(x) = 0$ for $x < 0$. The mean is unity, therefore $\zeta^{(w)}/\mathbf{E}(\zeta^{(w)})$ is an unbiased estimator of the mean, where $\mathbf{E}(\zeta^{(w)})$ is the expectation value of $\zeta^{(w)}$:

$$\mathbf{E}\,(\zeta^{(w)}) = \int_{-\infty}^{\infty} x g_w(x) dx. \tag{A4}$$

We can explicitly calculate $\mathbf{E}(\zeta^{(w)})$ for this case and the answer turns out to be given by a truncated alternating harmonic series:

$$\mathbf{E}\,(\zeta^{(w)}) = \sum_{j=1}^{w} \frac{(-1)^{j+1}}{j}. \tag{A5}$$

For very large w, $\mathbf{E}(\zeta^{(w)})$ approaches ln(2), which is precisely the population median. For $w = 1$ it is just unity, which makes sense because in this case the median is equal to the mean. For finite w, $\mathbf{E}(\zeta^{(w)})$ is somewhat larger than ln(2) and for a window size of $w = 101$, which is what is used in the actual search, it is 0.698 073. This is to be compared with ln(2) = 0.693 147, a difference of about 0.7%.

When w is even, the distribution of $\zeta^{(w)}$ is given (up to a factor of 2) by the convolution of the distributions of x_k and x_{k+1}. However, we are interested only in the expectation value

$$\mathbf{E}\,(\zeta^{(w)}) = \frac{1}{2}[\mathbf{E}(x_k) + \mathbf{E}(x_{k+1})]. \tag{A6}$$

The expectation value of x_{k+1} is calculated as above, using the distribution (A3), while the distribution of x_k is obtained by replacing k with $k - 1$ in (A3). It turns out that $\mathbf{E}(\zeta^{(w)}) = \mathbf{E}(\zeta^{(w-1)})$ when w is even. Thus $\mathbf{E}(\zeta^{(2)}) = \mathbf{E}(\zeta^{(1)})$, $\mathbf{E}(\zeta^{(4)}) = \mathbf{E}(\zeta^{(3)})$ and so on. Figure 18 plots $\mathbf{E}(\zeta^{(w)})$ for all values of w from 1 to 200.

APPENDIX B: THE NUMBER-COUNT OUTLIERS

This appendix contains a discussion about the outliers in the Hough number counts that are strongly suspected to be instrumental artifacts but that we were not able to definitely identify as such.

After applying the frequency veto described in Sec. VI D, we focus our attention on those candidates with a false alarm rate α (for a single detector) smaller than 10^{-13}. This corresponds to a threshold on the number count of 216 for L1, 480 for H1 and 390 for H2. Since the total number of templates analyzed in this 200 Hz search band is roughly 10^{12}, the probability of getting one candidate above that threshold over the full search is approximately 10% in each detector. All the candidates that satisfy such condition tend to cluster around a few frequencies that are listed in Table VI. These are the so-called *outliers* that were present in Figs. 5–7.

If these outliers are due to gravitational signals they should show up in the different detectors. By performing a simple coincidence analysis in frequency, the only candidate that remains, at this false alarm level, is the one at 210.36 Hz. The reader should notice that 210.36 Hz corresponds to 3×70.12 Hz. In the H1 data, there are also excess of number counts at 280.480 Hz and 350.600 Hz, corresponding to 4×70.12 and 5×70.12 respectively. These 70.120 Hz multiples together with the 244.14 Hz line were detected in association with a VME (VERSA module Eurocard) controller hardware used during S2. However, since the data acquisition system architecture has changed since S2 the coupling mechanism cannot be proven.

Figures 19–21 show how the outliers listed in Table VI stand well above the background noise spectrum level S_n,

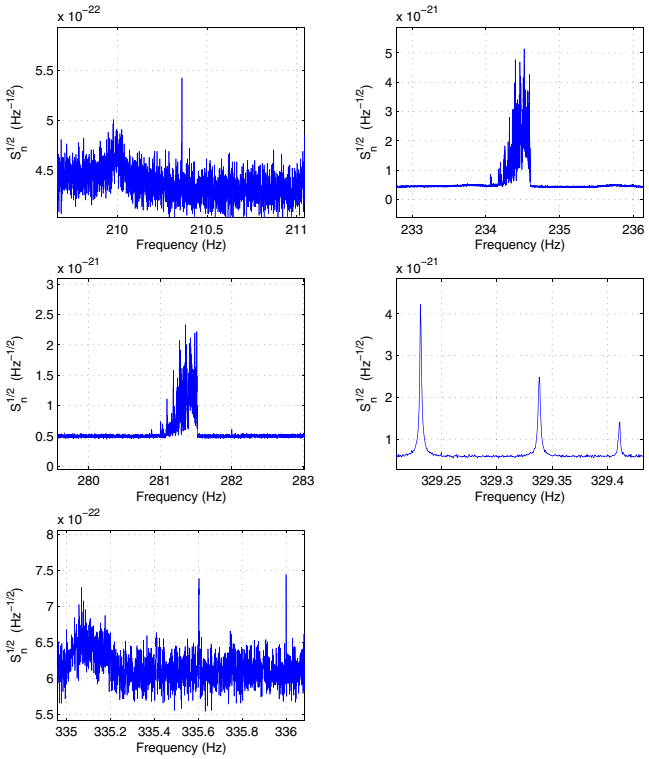

FIG. 19 (color online). The square root of the average value of S_n using the entire S2 L1 data set analyzed. The four graphs correspond to the frequencies where outliers were present in the Hough maps after applying the known instrumental veto. They correspond to 210.36 Hz, 234.50 Hz, 281.35 Hz, 329.34 Hz and 335.62 Hz.

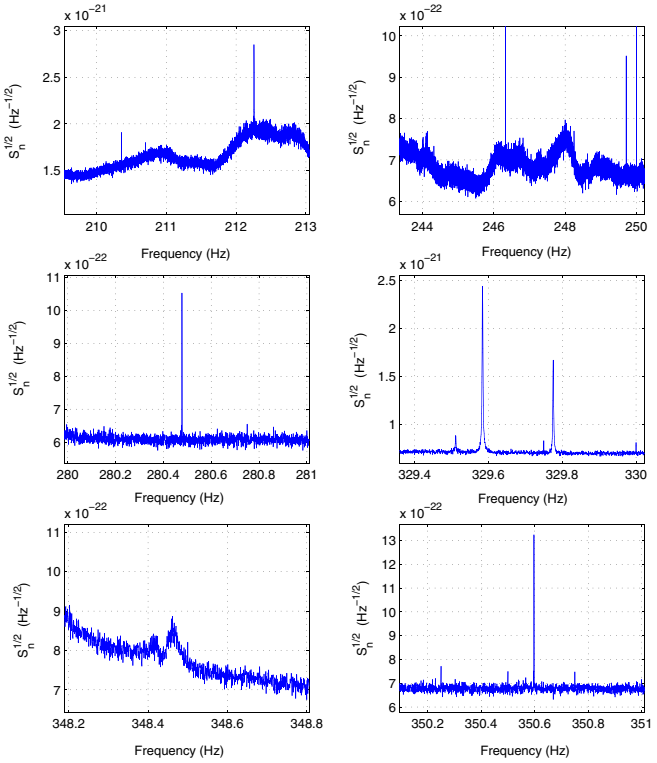

FIG. 20 (color online). The square root of the average value of S_n using the entire S2 H1 data set analyzed. The graphs correspond to zooms in the frequencies where outliers were present in the Hough maps. They correspond to 210.36 Hz, 212.26 Hz, 244.14 Hz, 249.70 Hz, 280.48 Hz, 329.58 Hz, 329.78 Hz, 348.45 Hz and 350.60 Hz.

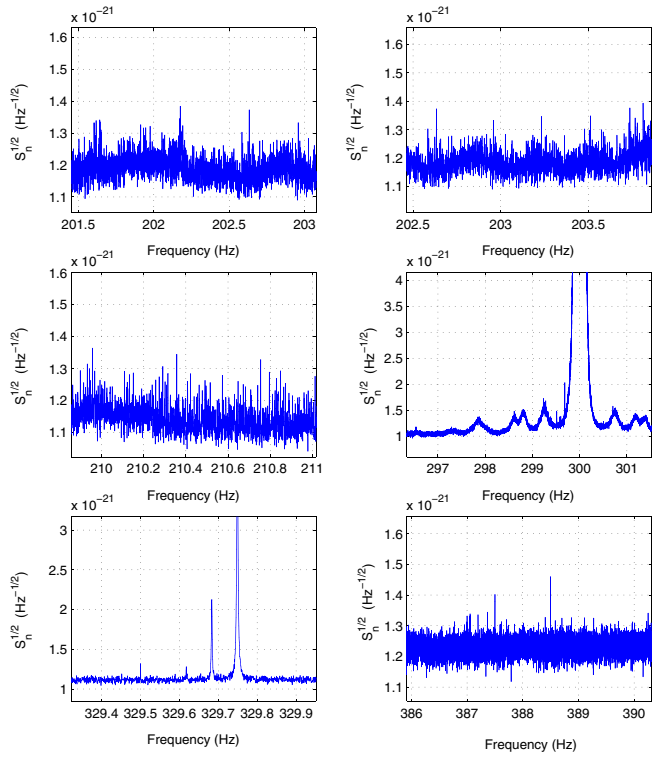

FIG. 21 (color online). The square root of the average value of S_n using the entire S2 H2 data set analyzed. The graphs correspond to zooms in the frequencies where outliers were present in the Hough maps after applying the known instrumental veto. Those correspond to 202.18 Hz, 203.23 Hz, 210.36 Hz, 298.81 Hz, 329.69 Hz, 387.05 Hz and 389.40 Hz.

when this is estimated from the entire run. We believe they all arise from instrumental or environmental artifacts. However we are not able to determine in a conclusive manner their physical cause. For this reason, the potential contaminated frequency bands have not been vetoed when setting the upper limits.

[1] B. Willke *et al.*, Classical Quantum Gravity **19**, 1377 (2002).
[2] S. Goßler *et al.*, Classical Quantum Gravity **19**, 1835 (2002).
[3] A. Abramovici *et al.*, Science **256**, 325 (1992).
[4] B. Barish and R. Weiss, Phys. Today **52** (10), 44 (1999).
[5] K. Tsubono, in *Proceedings of the 1st Edoardo Amaldi Conference on Gravitational Wave Experiments*, edited by E. Coccia, G. Pizella, and F. Ronga (World Scientific, Singapore, 1995), p. 112.
[6] B. Caron *et al.*, Nucl. Phys. B, Proc. Suppl. **54**, 167 (1997).
[7] L. Bildsten, Astrophys. J. **501**, L89 (1998).
[8] G. Ushomirsky, C. Cutler, and L. Bildsten, Mon. Not. R. Astron. Soc. **319**, 902 (2000).
[9] C. Cutler, Phys. Rev. D **66**, 084025 (2002).
[10] B. J. Owen, astro-ph/0503399 [Phys. Rev. Lett. (to be published)].
[11] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson, Phys. Rev. D **58**, 084020 (1998).
[12] N. Andersson, K. D. Kokkotas, and N. Stergioulas, Astrophys. J. **516**, 307 (1999).
[13] D. I. Jones and N. Andersson, Mon. Not. R. Astron. Soc. **331**, 203 (2002).
[14] C. Van Den Broeck, Classical Quantum Gravity **22**, 1825 (2005).
[15] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **69**, 082004 (2004).
[16] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. Lett. **94**, 181103 (2005).
[17] B. Abbott *et al.* (LIGO Scientific Collaboration), Nucl. Instrum. Methods Phys. Res., Sect. A **517**, 154 (2004).
[18] P. Astone, K. M. Borkowski, P. Jaranowski, and A. Królak,

Phys. Rev. D **65**, 042003 (2002).
[19] B. Abbott *et al.* (LIGO Scientific Collaboration), Report No. LIGO-P050008-00-Z, 2005.
[20] P. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Phys. Rev. D **57**, 2101 (1998).
[21] P. Brady and T. Creighton, Phys. Rev. D **61**, 082001 (2000).
[22] C. Cutler, I. Gholami, and B. Krishnan, Phys. Rev. D **72**, 042004 (2005).
[23] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, Phys. Rev. D **70**, 082001 (2004).
[24] M. A. Papa, B. F. Schutz, and A. M. Sintes, in *Gravitational Waves: A Challenge to Theoretical Astrophysics*, ICTP Lecture Notes Series Vol. III, edited by V. Ferrari, J. C. Miller, L. Rezzolla (Italy, 2001) p. 431.
[25] P. V. C. Hough, in *Proceedings of the International Conference on High Energy Accelerators and Instrumentation*, edited by L. Kowarski (CERN, Geneva, 1959) pp. 554–556.
[26] P. V. C. Hough, "Method and Means for Recognizing Complex Patterns," U.S. Patent No. 3 069 654, 1962.
[27] J. Illingworth and J. Kittler, Computer Vision, Graphics, and Image Processing **44**, 87 (1988).
[28] B. Krishnan *et al.* (LIGO Scientific Collaboration), Classical Quantum Gravity **22**, S1265 (2005).
[29] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **72**, 042002 (2005).
[30] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **72**, 062001 (2005).
[31] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **72**, 082001 (2005).
[32] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **72**, 082002 (2005).
[33] TAMA and LIGO Scientific Collaborations, gr-qc/0507081.
[34] B. Abbott *et al.* (LIGO Scientific Collaboration), gr-qc/0509129.
[35] B. Abbott *et al.* (LIGO Scientific Collaboration), astro-ph-qc/0507254 [Phys. Rev. Lett. (to be published)].
[36] G. González, M. Landry, B. O'Reilly, and H. Radkins, LIGO Technical Document No. T040060-01-D, 2004; available in http://admdbsrv.ligo.caltech.edu/dcc/
[37] D. R. Lorimer, Living Rev. Relativity **4**, 5 (2001); also at http://www.livingreviews.org/lrr-2001-5
[38] The Australian Telescope National Facility pulsar database: http://www.atnf.csiro.au/research/pulsar/psrcat/
[39] F. Haberl, Adv. Space Res. **33**, 638 (2004).
[40] G. G. Pavlov, D. Sanwal, and M. A. Teter, in *Young Neutron Stars and Their Environments, IAU Symposium 218*, ASP Conf. Proc., edited by F. Camilo and B. M. Gaensler (ASP, San Francisco, 2004), p. 239.
[41] V. M. Kaspi, M. S. E. Roberts, and A. K. Harding, astro-ph/0402136, edited by W. H. G. Lewin and M. van der Klis.
[42] C. Palomba, Mon. Not. R. Astron. Soc. **359**, 1150 (2005).
[43] J. Creighton, Phys. Rev. D **60**, 022001 (1999).
[44] B. Allen, M. A. Papa, and B. F. Schutz, Phys. Rev. D **66**, 102003 (2002).
[45] $T_{\rm obs}$ is the time from the first data to the last data of the SFTs used from the S2 run. Because the data are not continuous $T_{\rm obs} > NT_{\rm coh}$ and it is also different for the three detectors.
[46] The spin-down resolution depends on the total observation time, and this turns out to be $-1.105\,08 \times 10^{-10}$ Hz-s^{-1} for L1, $-1.091\,59 \times 10^{-10}$ Hz-s^{-1} for H1, and $-1.103\,88 \times 10^{-10}$Hz-s^{-1} for H2.
[47] The window size w should be as small as possible in order to track the noise floor accurately. On the other hand, the statistical errors in the value of the running median are smaller when a large window size is chosen. We have found that $w = 101$ is a good compromise.
[48] S. D. Mohanty, Classical Quantum Gravity **19**, 1513 (2002).
[49] S. D. Mohanty and S. Mukherjee, Classical Quantum Gravity **19**, 1471 (2002).
[50] B. Krishnam, LIGO Technical Document No. T-040144-00-Z, 2004; available in http://admdbsrv.ligo.caltech.edu/dcc/
[51] L. Brocco, S. Frasca, C. Palomba, and F. Ricci, Classical Quantum Gravity **20**, S655 (2003).
[52] http://www.cs.wisc.edu/condor/
[53] http://pandora.aei.mpg.de/merlin/
[54] http://www.lsc-group.phys.uwm.edu/cgi-bin/cvs/ viewcvs.cgi/?cvsroot=lscsoft
[55] M. Landry, LIGO Technical Document No. T050089-00-D, 2005; available in http://admdbsrv.ligo.caltech.edu/dcc/
[56] R. Prix and Y. Itoh, Classical Quantum Gravity **22**, S1003 (2005).
[57] P. Jaranowski, A. Królak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998).
[58] H. Kramer, *Mathematical Methods of Statistics, Princeton Landmarks in Mathematics* (Princeton University Press, Princeton, NJ, 1946).