Searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: Results from the second LIGO science run

B. Abbott,[15] R. Abbott,[15] R. Adhikari,[15] J. Agresti,[15] P. Ajith,[2] B. Allen,[2,54] R. Amin,[19] S. B. Anderson,[15] W. G. Anderson,[54] M. Arain,[41] M. Araya,[15] H. Armandula,[15] M. Ashley,[4] S Aston,[40] P. Aufmuth,[38] C. Aulbert,[1] S. Babak,[1] S. Ballmer,[15] H. Bantilan,[9] B. C. Barish,[15] C. Barker,[16] D. Barker,[16] B. Barr,[42] P. Barriga,[53] M. A. Barton,[42] K. Bayer,[18] K. Belczynski,[26] S. J. Berukoff,[1] J. Betzwieser,[18] P. T. Beyersdorf,[29] B. Bhawal,[15] I. A. Bilenko,[23] G. Billingsley,[15] R. Biswas,[54] E. Black,[15] K. Blackburn,[15] L. Blackburn,[18] B. Blair,[53] B. Bland,[16] J. Bogenstahl,[42] L. Bogue,[17] R. Bork,[15] V. Boschi,[15] S. Bose,[56] P. R. Brady,[54] V. B. Braginsky,[23] J. E. Brau,[45] M. Brinkmann,[2] A. Brooks,[39] D. A. Brown,[15,7] A. Bullington,[32] A. Bunkowski,[2] A. Buonanno,[43] O. Burmeister,[2] D. Busby,[15] W. E. Butler,[46] R. L. Byer,[32] L. Cadonati,[18] G. Cagnoli,[42] J. B. Camp,[24] J. Cannizzo,[24] K. Cannon,[54] C. A. Cantley,[42] J. Cao,[18] L. Cardenas,[15] K. Carter,[17] M. M. Casey,[42] G. Castaldi,[48] C. Cepeda,[15] E. Chalkey,[42] P. Charlton,[10] S. Chatterji,[15] S. Chelkowski,[2] Y. Chen,[1] F. Chiadini,[47] D. Chin,[44] E. Chin,[53] J. Chow,[4] N. Christensen,[9] J. Clark,[42] .P Cochrane,[2] T. Cokelaer,[8] C. N. Colacino,[40] R. Coldwell,[41] M. Coles,[17] R. Conte,[47] D. Cook,[16] T. Corbitt,[18] D. Coward,[53] D. Coyne,[15] J. D. E. Creighton,[54] T. D. Creighton,[15] R. P. Croce,[48] D. R. M. Crooks,[42] A. M. Cruise,[40] P. Csatorday,[18] A. Cumming,[42] C. Cutler,[7] J. Dalrymple,[33] E. D'Ambrosio,[15] K. Danzmann,[38,2] G. Davies,[8] E. Daw,[49] D. DeBra,[32] J. Degallaix,[53] M. Degree,[32] T. Delker,[41] T. Demma,[48] V. Dergachev,[44] S. Desai,[34] R. DeSalvo,[15] S. Dhurandhar,[14] M. Díaz,[35] J. Dickson,[4] A. Di Credico,[33] G. Diederichs,[38] A. Dietz,[8] H. Ding,[15] E. E. Doomes,[31] R. W. P. Drever,[5] J.-C. Dumas,[53] R. J. Dupuis,[15] J. G. Dwyer,[11] P. Ehrens,[15] E. Espinoza,[15] T. Etzel,[15] M. Evans,[15] T. Evans,[17] S. Fairhurst,[8,15] Y. Fan,[53] D. Fazi,[15] M. M. Fejer,[32] L. S. Finn,[34] V. Fiumara,[47] N. Fotopoulos,[54] A. Franzen,[38] K. Y. Franzen,[41] A. Freise,[40] R. Frey,[45] T. Fricke,[46] P. Fritschel,[18] V. V. Frolov,[17] M. Fyffe,[17] V. Galdi,[48] K. S. Ganezer,[6] J. Garofoli,[16] I. Gholami,[1] J. A. Giaime,[17,19] S. Giampanis,[46] K. D. Giardina,[17] K. Goda,[18] E. Goetz,[44] L. M. Goggin,[15] G. González,[19] S. Gossler,[4] A. Grant,[42] S. Gras,[53] C. Gray,[16] M. Gray,[4] J. Greenhalgh,[28] A. M. Gretarsson,[12] R. Grosso,[35] H. Grote,[2] S. Grunewald,[1] M. Guenther,[16] R. Gustafson,[44] B. Hage,[38] D. Hammer,[54] C. Hanna,[19] J. Hanson,[17] J. Harms,[2] G. Harry,[18] E. Harstad,[45] T. Hayler,[28] J. Heefner,[15] G. Heinzel,[2] I. S. Heng,[42] A. Heptonstall,[42] M. Heurs,[2] M. Hewitson,[2] S. Hild,[38] E. Hirose,[33] D. Hoak,[17] D. Hosken,[39] J. Hough,[42] E. Howell,[53] D. Hoyland,[40] S. H. Huttner,[42] D. Ingram,[16] E. Innerhofer,[18] M. Ito,[45] Y. Itoh,[54] A. Ivanov,[15] D. Jackrel,[32] O. Jennrich,[42] B. Johnson,[16] W. W. Johnson,[19] W. R. Johnston,[35] D. I. Jones,[50] G. Jones,[8] R. Jones,[42] L. Ju,[53] P. Kalmus,[11] V. Kalogera,[26] D. Kasprzyk,[40] E. Katsavounidis,[18] K. Kawabe,[16] S. Kawamura,[25] F. Kawazoe,[25] W. Kells,[15] D. G. Keppel,[15] F. Ya. Khalili,[23] C. J. Killow,[42] C. Kim,[26] P. King,[15] J. S. Kissell,[19] S. Klimenko,[41] K. Kokeyama,[25] V. Kondrashov,[15] R. K. Kopparapu,[19] D. Kozak,[15] B. Krishnan,[1] P. Kwee,[38] P. K. Lam,[4] M. Landry,[16] B. Lantz,[32] A. Lazzarini,[15] B. Lee,[53] M. Lei,[15] J. Leiner,[56] V. Leonhardt,[25] I. Leonor,[45] K. Libbrecht,[15] A. Libson,[9] P. Lindquist,[15] N. A. Lockerbie,[51] J. Logan,[15] M. Longo,[47] M. Lormand,[17] M. Lubinski,[16] H. Lück,[38,2] B. Machenschalk,[1] M. MacInnis,[18] M. Mageswaran,[15] K. Mailand,[15] M. Malec,[38] V. Mandic,[15] S. Marano,[47] S. Márka,[11] J. Markowitz,[18] E. Maros,[15] I. Martin,[42] J. N. Marx,[15] K. Mason,[18] L. Matone,[11] V. Matta,[47] N. Mavalvala,[18] R. McCarthy,[16] D. E. McClelland,[4] S. C. McGuire,[31] M. McHugh,[21] K. McKenzie,[4] J. W. C. McNabb,[34] S. McWilliams,[24] T. Meier,[38] A. Melissinos,[46] G. Mendell,[16] R. A. Mercer,[41] S. Meshkov,[15] E. Messaritaki,[15] C. J. Messenger,[42] D. Meyers,[15] E. Mikhailov,[18] S. Mitra,[14] V. P. Mitrofanov,[23] G. Mitselmakher,[41] R. Mittleman,[18] O. Miyakawa,[15] S. Mohanty,[35] G. Moreno,[16] K. Mossavi,[2] C. MowLowry,[4] A. Moylan,[4] D. Mudge,[39] G. Mueller,[41] S. Mukherjee,[35] H. Müller-Ebhardt,[2] J. Munch,[39] P. Murray,[42] E. Myers,[16] J. Myers,[16] S. Nagano,[2] T. Nash,[15] G. Newton,[42] A. Nishizawa,[25] F. Nocera,[15] K. Numata,[24] P. Nutzman,[26] B. O'Reilly,[17] R. O'Shaughnessy,[26] D. J. Ottaway,[18] H. Overmier,[17] B. J. Owen,[34] Y. Pan,[43] M. A. Papa,[1,54] V. Parameshwaraiah,[16] C. Parameswariah,[17] P. Patel,[15] M. Pedraza,[15] S. Penn,[13] V. Pierro,[48] I. M. Pinto,[48] M. Pitkin,[42] H. Pletsch,[2] M. V. Plissi,[42] F. Postiglione,[47] R. Prix,[1] V. Quetschke,[41] F. Raab,[16] D. Rabeling,[4] H. Radkins,[16] R. Rahkola,[45] N. Rainer,[2] M. Rakhmanov,[34] M. Ramsunder,[34] K. Rawlins,[18] S. Ray-Majumder,[54] V. Re,[40] T. Regimbau,[8] H. Rehbein,[2] S. Reid,[42] D. H. Reitze,[41] L. Ribichini,[2] S. Richman,[18] R. Riesen,[17] K. Riles,[44] B. Rivera,[16] N. A. Robertson,[15,42] C. Robinson,[8] E. L. Robison,[40] S. Roddy,[17] A. Rodriguez,[19] A. M. Rogan,[56] J. Rollins,[11] J. D. Romano,[8] J. Romie,[17] H. Rong,[41] R. Route,[32] S. Rowan,[42] A. Rüdiger,[2] L. Ruet,[18] P. Russell,[15] K. Ryan,[16] S. Sakata,[25] M. Samidi,[15] L. Sancho de la Jordana,[37] V. Sandberg,[16] G. H. Sanders,[15] V. Sannibale,[15] S. Saraf,[27] P. Sarin,[18] B. Sathyaprakash,[8] S. Sato,[25] P. R. Saulson,[33] R. Savage,[16] P. Savov,[7] A. Sazonov,[41] S. Schediwy,[53] R. Schilling,[2] R. Schnabel,[2] R. Schofield,[45] B. F. Schutz,[1,8] P. Schwinberg,[16] S. M. Scott,[4] A. C. Searle,[4] B. Sears,[15] F. Seifert,[2] D. Sellers,[17] A. S. Sengupta,[8] P. Shawhan,[43] D. H. Shoemaker,[18] A. Sibley,[17] J. A. Sidles,[52] X. Siemens,[15,7] D. Sigg,[16] S. Sinha,[32] A. M. Sintes,[37,1] B. J. J. Slagmolen,[4] J. Slutsky,[19] J. R. Smith,[2] M. R. Smith,[15] K. Somiya,[2,1] K. A. Strain,[42]

N. E. Strand,[34] D. M. Strom,[45] A. Stuver,[34] T. Z. Summerscales,[3] K.-X. Sun,[32] M. Sung,[19] P. J. Sutton,[15] J. Sylvestre,[15] H. Takahashi,[1] A. Takamori,[15] D. B. Tanner,[41] M. Tarallo,[15] R. Taylor,[15] R. Taylor,[42] J. Thacker,[17] K. A. Thorne,[34] K. S. Thorne,[7] A. Thüring,[38] M. Tinto,[5] K. V. Tokmakov,[42] C. Torres,[35] C. Torrie,[42] G. Traylor,[17] M. Trias,[37] W. Tyler,[15] D. Ugolini,[36] C. Ungarelli,[40] K. Urbanek,[32] H. Vahlbruch,[38] M. Vallisneri,[7] C. Van Den Broeck,[8] M. van Putten,[18] M. Varvella,[15] S. Vass,[15] A. Vecchio,[40] J. Veitch,[42] P. Veitch,[39] A. Villar,[15] C. Vorvick,[16] S. P. Vyachanin,[23] S. J. Waldman,[15] L. Wallace,[15] H. Ward,[42] R. Ward,[15] K. Watts,[17] D. Webber,[15] A. Weidner,[2] M. Weinert,[2] A. Weinstein,[15] R. Weiss,[18] L. Wen,[1] S. Wen,[19] K. Wette,[4] J. T. Whelan,[1] D.. M. Whitbeck,[34] S. E. Whitcomb,[15] B. F. Whiting,[41] S. Wiley,[6] C. Wilkinson,[16] P. A. Willems,[15] L. Williams,[41] B. Willke,[38,2] I. Wilmut,[28] W. Winkler,[2] C. C. Wipf,[18] S. Wise,[41] A. G. Wiseman,[54] G. Woan,[42] D. Woods,[54] R. Wooley,[17] J. Worden,[16] W. Wu,[41] I. Yakushin,[17] H. Yamamoto,[15] Z. Yan,[53] S. Yoshida,[30] N. Yunes,[34] K. D. Zaleski,[34] M. Zanolin,[18] J. Zhang,[44] L. Zhang,[15] C. Zhao,[53] N. Zotov,[20] M. Zucker,[18] H. zur Mühlen,[38] and J. Zweizig[15]

(LIGO Scientific Collaboration)*

[1]*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476 Golm, Germany*
[2]*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany*
[3]*Andrews University, Berrien Springs, Michigan 49104, USA*
[4]*Australian National University, Canberra, 0200, Australia*
[5]*California Institute of Technology, Pasadena, California 91125, USA*
[6]*California State University, Dominguez Hills, Carson, California 90747, USA*
[7]*Caltech-CaRT, Pasadena, California 91125, USA*
[8]*Cardiff University, Cardiff, CF2 3YB, United Kingdom*
[9]*Carleton College, Northfield, Minnesota 55057, USA*
[10]*Charles Sturt University, Wagga Wagga, NSW 2678, Australia*
[11]*Columbia University, New York, New York 10027, USA*
[12]*Embry-Riddle Aeronautical University, Prescott, AZ 86301 USA*
[13]*Hobart and William Smith Colleges, Geneva, New York 14456, USA*
[14]*Inter-University Centre for Astronomy and Astrophysics, Pune - 411007, India*
[15]*LIGO - California Institute of Technology, Pasadena, California 91125, USA*
[16]*LIGO Hanford Observatory, Richland, Washington 99352, USA*
[17]*LIGO Livingston Observatory, Livingston, Louisiana 70754, USA*
[18]*LIGO - Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*
[19]*Louisiana State University, Baton Rouge, Louisiana 70803, USA*
[20]*Louisiana Tech University, Ruston, Louisiana 71272, USA*
[21]*Loyola University, New Orleans, Louisiana 70118, USA*
[22]*Max Planck Institut für Quantenoptik, D-85748, Garching, Germany*
[23]*Moscow State University, Moscow, 119992, Russia*
[24]*NASA/Goddard Space Flight Center, Greenbelt, Maryland 20771, USA*
[25]*National Astronomical Observatory of Japan, Tokyo 181-8588, Japan*
[26]*Northwestern University, Evanston, Illinois 60208, USA*
[27]*Rochester Institute of Technology, Rochester, NY 14623, USA*
[28]*Rutherford Appleton Laboratory, Chilton, Didcot, Oxon OX11 0QX United Kingdom*
[29]*San Jose State University, San Jose, California 95192, USA*
[30]*Southeastern Louisiana University, Hammond, Louisiana 70402, USA*
[31]*Southern University, Baton Rouge, LA 70813, USA and A&M College, Baton Rouge, LA 70813, USA*
[32]*Stanford University, Stanford, California 94305, USA*
[33]*Syracuse University, Syracuse, New York 13244, USA*
[34]*The Pennsylvania State University, University Park, Pennsylvania 16802, USA*
[35]*The University of Texas at Brownsville, Brownsville, Texas 78520, USA*
and Texas Southmost College, Brownsville, Texas 78520, USA
[36]*Trinity University, San Antonio, Texas 78212, USA*
[37]*Universitat de les Illes Balears, E-07122 Palma de Mallorca, Spain*
[38]*Universität Hannover, D-30167 Hannover, Germany*
[39]*University of Adelaide, Adelaide, SA 5005, Australia*
[40]*University of Birmingham, Birmingham, B15 2TT, United Kingdom*
[41]*University of Florida, Gainesville, Florida 32611, USA*
[42]*University of Glasgow, Glasgow, G12 8QQ, United Kingdom*
[43]*University of Maryland, College Park, Massachusetts 20742, USA*

[44]University of Michigan, Ann Arbor, Michigan 48109, USA
[45]University of Oregon, Eugene, Oregon 97403, USA
[46]University of Rochester, Rochester, New York 14627, USA
[47]University of Salerno, 84084 Fisciano (Salerno), Italy
[48]University of Sannio at Benevento, I-82100 Benevento, Italy
[49]University of Sheffield, Sheffield, S3 7RH, United Kingdom
[50]University of Southampton, Southampton, SO17 1BJ, United Kingdom
[51]University of Strathclyde, Glasgow, G1 1XQ, United Kingdom
[52]University of Washington, Seattle, Washington 98195, USA
[53]University of Western Australia, Crawley, WA 6009, Australia
[54]University of Wisconsin-Milwaukee, Milwaukee, Wisconsin 53201, USA
[55]Vassar College, Poughkeepsie, New York 12604, USA
[56]Washington State University, Pullman, Washington 99164, USA
(Received 12 June 2006; revised manuscript received 2 April 2007; published 24 October 2007)

We carry out two searches for periodic gravitational waves using the most sensitive few hours of data from the second LIGO science run. Both searches exploit fully coherent matched filtering and cover wide areas of parameter space, an innovation over previous analyses which requires considerable algorithm development and computational power. The first search is targeted at isolated, previously unknown neutron stars, covers the entire sky in the frequency band 160–728.8 Hz, and assumes a frequency derivative of less than 4×10^{-10} Hz/s. The second search targets the accreting neutron star in the low-mass x-ray binary Scorpius X-1 and covers the frequency bands 464–484 Hz and 604–624 Hz as well as the two relevant binary orbit parameters. Because of the high computational cost of these searches we limit the analyses to the most sensitive 10 hours and 6 hours of data, respectively. Given the limited sensitivity and duration of the analyzed data set, we do not attempt deep follow-up studies. Rather we concentrate on demonstrating the data analysis method on a real data set and present our results as upper limits over large volumes of the parameter space. In order to achieve this, we look for coincidences in parameter space between the Livingston and Hanford 4-km interferometers. For isolated neutron stars our 95% confidence level upper limits on the gravitational wave strain amplitude range from 6.6×10^{-23} to 1×10^{-21} across the frequency band; for Scorpius X-1 they range from 1.7×10^{-22} to 1.3×10^{-21} across the two 20-Hz frequency bands. The upper limits presented in this paper are the first broadband wide parameter space upper limits on periodic gravitational waves from coherent search techniques. The methods developed here lay the foundations for upcoming hierarchical searches of more sensitive data which may detect astrophysical signals.

DOI: 10.1103/PhysRevD.76.082001 PACS numbers: 04.80.Nn, 07.05.Kf, 95.55.Ym, 97.60.Gb

I. INTRODUCTION

Rapidly rotating neutron stars are the most likely sources of persistent gravitational radiation in the frequency band $\approx$ 100 Hz–1 kHz. These objects may generate continuous gravitational waves (GW) through a variety of mechanisms, including nonaxisymmetric distortions of the star [1–5], velocity perturbations in the star's fluid [1,6,7], and free precession [8,9]. Regardless of the specific mechanism, the emitted signal is a quasiperiodic wave whose frequency changes slowly during the observation time due to energy loss through gravitational wave emission, and possibly other mechanisms. At an Earth-based detector the signal exhibits amplitude and phase modulations due to the motion of the Earth with respect to the source. The intrinsic gravitational wave amplitude is likely to be several orders of magnitude smaller than the typical root-mean-square value of the detector noise, hence detection can only be achieved by means of long integration times, of the order of weeks to months.

Deep, wide parameter space searches for continuous gravitational wave signals are computationally bound. At fixed computational resources the optimal sensitivity is achieved through hierarchical search schemes [10–12]. Such schemes alternate incoherent and coherent search stages in order to first efficiently identify statistically significant candidates and then follow them up with more sensitive, albeit computationally intensive, methods. Hierarchical search schemes have been investigated only theoretically, under the simplified assumption of Gaussian and stationary instrumental noise; the computational costs have been estimated only on the basis of rough counts of floating point operations necessary to evaluate some detection statistic, usually not the optimal, and have not taken into account additional costs coming e.g. from data input/output; computational savings obtainable through efficient dedicated numerical implementations have also been neglected. Furthermore, general theoretical investigations

*Electronic address: http://www.ligo.org

have not relied on the optimizations that can be introduced on the basis of the specific area in parameter space at which a search is aimed.

In this paper we demonstrate and characterize the coherent stage of a hierarchical pipeline by carrying out two large parameter space coherent searches on data collected by LIGO during the second science run with the Livingston and Hanford 4-km interferometers. As we will show, this analysis requires careful tuning of a variety of search parameters and implementation choices, such as the tilings of the parameter space, and the selection of the data that are difficult to determine on purely theoretical grounds. This paper complements the study presented in [13] where we reported results obtained by applying an incoherent analysis method [14] to data taken during the same science run. Furthermore, here we place upper limits on regions of the parameter space that have never been explored before. We do this by combining the output of the coherent searches via a coincidence scheme.

The coherent search described in this paper has been the test-bench for the core science analysis that the Einstein@home [15] project is carrying out now. The development of analysis techniques such as the one described here, together with the computing power of Einstein@home in the context of a hierarchical search scheme, will allow the deepest searches for continuous gravitational waves.

In this paper the same basic pipeline is applied to and tuned for two different searches: (i) for signals from isolated sources over the whole sky and the frequency band 160 Hz–728.8 Hz, and (ii) for a signal from the low-mass x-ray binary Scorpius X-1 (Sco X-1) over orbital parameters and in the frequency bands 464–484 Hz and 604–624 Hz. It is the first time that a coherent analysis is carried out over such a wide frequency band and coincidence techniques are used among the registered candidates from different detectors; the only other example of a somewhat similar analysis is an all-sky search over two days of data from the Explorer resonant detector and that was over a 0.76 Hz band around 922 Hz [16–18]. This is absolutely the first wide parameter space search for a rotating neutron star in a binary system.

The main scope of the paper is to illustrate an analysis method by applying it to two different wide parameter spaces. In fact, based on the typical noise performance of the detectors during the run, which is shown in Fig. 1, and the amount of data that we were able to process in $\approx$ 1 month with our computational resources (totalling about 800 CPUs over several Beowulf clusters) we do not expect to detect gravitational waves. For isolated neutron stars we estimate (see Sec. III for details) that statistically the strongest signal that we expect from an isolated source is $\lesssim 4 \times 10^{-24}$ which is a factor $\gtrsim 20$ smaller than the dimmest signal that we would have been able to observe with the present search. For Scorpius X-1, the signal is

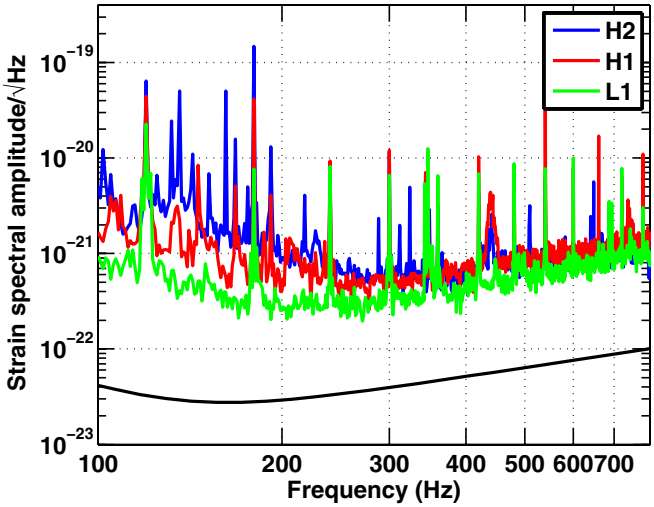

FIG. 1 (color online). Typical one-sided amplitude spectral densities of detector noise during the second science run, for the three LIGO instruments. The solid black line is the design sensitivity for the two 4-km instruments L1 and H1.

expected to have a strength of at most $\sim 3 \times 10^{-26}$ and our search is a factor $\sim$5000 less sensitive. The results of the analyses confirm these expectations and we report upper limits for both searches.

The paper is organized as follows. In Sec. II we describe the instrument configuration during the second science run and the details of the data taking. In Sec. III we review the current astrophysical understanding of neutron stars as gravitational wave sources, including a somewhat novel statistical argument that the strength of the strongest such signal that we can expect to receive does not exceed $h_0^{\max} \approx 4 \times 10^{-24}$. We also detail and motivate the choice of parameter spaces explored in this paper. In Sec. IV we review the signal model and discuss the search area considered here. In Sec. V we describe the analysis pipeline. In Sec. VI we present and discuss the results of the analyses. In Sec. VII we recapitulate the most relevant results in the wider context and provide pointers for future work.

II. INSTRUMENTS AND THE SECOND SCIENCE RUN

Three detectors at two independent sites comprise the Laser Interferometer Gravitational Wave Observatory, or LIGO. Detector commissioning has progressed since the fall of 1999, interleaved with periods in which the observatory ran nearly continuously for weeks or months, the so-called ''science runs.'' The first science run (S1) was made in concert with the gravitational wave detector GEO600; results from the analysis of those data were presented in [19–22], while the instrument status was detailed in [23]. Significant improvements in the strain sensitivity of the LIGO interferometers (an order of magnitude over a broadband) culminated in the second science run (S2), which

took place from February 14 to April 14, 2003. Details of the S2 run, including detector improvements between S1 and S2 can be found in [24], Sec. IV of [25], and Sec. II of [13,26].

Each LIGO detector is a recycled Michelson interferometer with Fabry-Perot arms, whose lengths are defined by suspended mirrors that double as test masses. Two detectors reside in the same vacuum in Hanford, WA, one (denoted H1) with 4-km armlength and one with 2-km armlength (H2), while a single 4-km counterpart (L1) exists in Livingston Parish, LA. Differential motions are sensed interferometrically, and the resultant sensitivity is broadband (40 Hz–7 kHz), with spectral disturbances such as 60 Hz powerline harmonics evident in the noise spectrum (see Fig. 1). Optical resonance, or "lock," in a given detector is maintained by servo loops; lock may be interrupted by, for example, seismic transients or poorly conditioned servos. S2 duty cycles, accounting for periods in which lock was broken and/or detectors were known to be functioning not at the required level, were 74% for H1, 58% for H2, and 37% for L1. The two analyses described in this paper used a small subset of the data from the two most sensitive instruments during S2, L1, and H1; the choice of the segments considered for the analysis is detailed in Sec. V C.

The strain signal at the interferometer output is reconstructed from the error signal of the feedback loop which is used to control the differential length of the arms of the instrument. Such a process—known as calibration—involves the injection of continuous, constant amplitude sinusoidal excitations into the end test mass control systems, which are then monitored at the measurement error point. The calibration process introduces uncertainties in the amplitude of the recorded signal that were estimated to be $\lesssim 11\%$ during S2 [27]. In addition, during the run artificial pulsar-like signals were injected into the data stream by physically moving the mirrors of the Fabry-Perot cavity. Such "hardware injections" were used to validate the data analysis pipeline and details are presented in Appendix C.

III. ASTROPHYSICAL SOURCES

We review the physical mechanisms of periodic gravitational wave emission and the target populations of the two searches described in this paper. We also compare the sensitivity of these searches to likely source strengths.

A. Emission mechanisms

In the LIGO frequency band there are three predicted mechanisms for producing periodic gravitational waves, all of which involve neutron stars or similar compact objects: (1) nonaxisymmetric distortions of the solid part of the star [1–5], (2) unstable r-modes in the fluid part of the star [1,6,7], and (3) free precession of the whole star [8,9].

We begin with nonaxisymmetric distortions. These could not exist in a perfect fluid star, but in realistic neutron stars such distortions could be supported either by elastic stresses or by magnetic fields. The deformation is often expressed in terms of the ellipticity

$$\epsilon = \frac{I_{xx} - I_{yy}}{I_{zz}}, \tag{1}$$

which is (up to a numerical factor of order unity) the $m = 2$ quadrupole moment divided by the principal moment of inertia. A nonaxisymmetric neutron star rotating with frequency ν emits periodic gravitational waves with amplitude

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

where G is Newton's gravitational constant, c is the speed of light, I_{zz} is the principal moment of inertia of the object, f (equal to 2ν) is the gravitational wave frequency, and d is the distance to the object. Equation (2) gives the strain amplitude of a gravitational wave from an optimally oriented source [see Eq. (25) below].

The ellipticity of neutron stars is highly uncertain. The maximum ellipticity that can be supported by a neutron star's crust is estimated to be [2]

$$\epsilon_{\max} \approx 5 \times 10^{-7}\left(\frac{\sigma}{10^{-2}}\right), \tag{3}$$

where σ is the breaking strain of the solid crust. The numerical coefficient in Eq. (3) is small mainly because the shear modulus of the inner crust (which constitutes most of the crust's mass) is small, in the sense that it is about 10^{-3} times the pressure. Equation (3) uses a fiducial breaking strain of 10^{-2} since that is roughly the upper limit for the best terrestrial alloys. However, σ could be as high as 10^{-1} for a perfect crystal with no defects [28], or several orders of magnitude smaller for an amorphous solid or a crystal with many defects.

Some exotic alternatives to standard neutron stars feature solid cores, which could support considerably larger ellipticities [5]. The most speculative and highest-ellipticity model is that of a solid strange-quark star, for which

$$\epsilon_{\max} \approx 4 \times 10^{-4}\left(\frac{\sigma}{10^{-2}}\right). \tag{4}$$

This much higher value of $\epsilon_{\max}$ is mostly due to the higher shear modulus, which for some strange star models can be almost as large as the pressure. Another (still speculative but more robust) model is the hybrid star, which consists of a normal neutron star outside a solid core of mixed quark and baryon matter, which may extend from the center to nearly the bottom of the crust. For hybrid stars,

$$\epsilon_{\max} \approx 9 \times 10^{-6}\left(\frac{\sigma}{10^{-2}}\right), \tag{5}$$

although this is highly dependent on the poorly known range of densities occupied by the quark-baryon mixture. Stars with charged meson condensates could also have solid cores with overall ellipticities similar to those of hybrid stars.

Regardless of the maximum ellipticity supportable by shear stresses, there is the separate problem of how to reach the maximum. The crust of a young neutron star probably cracks as the neutron star spins down, but it is unclear how long it takes for gravity to smooth out the neutron star's shape. Accreting neutron stars in binaries have a natural way of reaching and maintaining the maximum deformation, since the accretion flow, guided by the neutron star's magnetic field, naturally produces "hot spots" on the surface, which can imprint themselves as lateral temperature variations throughout the crust. Through the temperature dependence of electron capture, these variations can lead to "hills" in hotter areas which extend down to the dense inner crust, and with a reasonable temperature variation the ellipticity might reach the maximum elastic value [1]. The accreted material can also be held up in mountains on the surface by the magnetic field itself: The matter is a good conductor, and thus it crosses field lines relatively slowly and can pile up in mountains larger than those supportable by elasticity alone [4]. Depending on the field configuration, accretion rate, and temperature, the ellipticity from this mechanism could be up to 10^{-5} even for ordinary neutron stars.

Strong internal magnetic fields are another possible cause of ellipticity [3]. Differential rotation immediately after the core collapse in which a neutron star is formed can lead to an internal magnetic field with a large toroidal part. Dissipation tends to drive the symmetry axis of a toroidal field toward the star's equator, which is the orientation that maximizes the ellipticity. The resulting ellipticity is

$$\epsilon \approx \begin{cases} 1.6 \times 10^{-6}\left(\frac{B}{10^{15}\ \mathrm{G}}\right) & B < 10^{15}\ \mathrm{G}, \\ 1.6 \times 10^{-6}\left(\frac{B}{10^{15}\ \mathrm{G}}\right)^2 & B > 10^{15}\ \mathrm{G}, \end{cases} \tag{6}$$

where B is the root-mean-square value of the toroidal part of the field averaged over the interior of the star. Note that this mechanism requires that the external field be much smaller than the internal field, since such strong external fields will spin a star out of the LIGO frequency band on a very short time scale.

An alternative way of generating asymmetry is the r-modes, fluid oscillations dominated by the Coriolis restoring force. These modes may be unstable to growth through gravitational radiation reaction [the Chandrasekhar-Friedman-Schutz (CFS) instability] under astrophysically realistic conditions. Rather than go into the many details of the physics and astrophysics, we refer the reader to a recent review [29] of the literature and summarize here only what is directly relevant to our search: The r-modes have been proposed as a source of gravitational waves from newborn neutron stars [6] and from rapidly accreting neutron stars [1,7]. The CFS instability of the r-modes in newborn neutron stars is probably not a good candidate for detection because the emission is very short-lived, low amplitude, or both. Accreting neutron stars (or quark stars) are a better prospect for a detection of r-mode gravitational radiation because the emission may be long-lived with a duty cycle near unity [30,31].

Finally we consider free precession, i.e. the wobble of a neutron star whose symmetry axis does not coincide with its rotation axis. A large-amplitude wobble would produce [8]

$$h_0 \sim 10^{-27}\left(\frac{\theta_{\mathrm{w}}}{0.1}\right)\left(\frac{1\ \mathrm{kpc}}{d}\right)\left(\frac{\nu}{500\ \mathrm{Hz}}\right)^2, \tag{7}$$

where θ_{w} is the wobble amplitude in radians. Such wobble may be longer lived than previously thought [9], but the amplitude is still small enough that such radiation is a target for second generation interferometers such as Advanced LIGO.

In light of our current understanding of emission mechanisms, the most likely sources of detectable gravitational waves are isolated neutron stars (through deformations) and accreting neutron stars in binaries (through deformations or r-modes).

B. Isolated neutron stars

The target population of this search is isolated rotating compact stars that have not been observed electromagnetically. Current models of stellar evolution suggest that our Galaxy contains of order 10^9 neutron stars, while only of order 10^5 are active pulsars. Up to now only about 1500 have been observed [32]; there are numerous reasons for this, including selection effects and the fact that many have faint emission. Therefore the target population is a large fraction of the neutron stars in the Galaxy.

1. Maximum expected signal amplitude at the earth

Despite this large target population and the variety of GW emission mechanisms that have been considered, one can make a robust argument, based on energetics and statistics, that the amplitude of the strongest gravitational wave pulsar that one could reasonably hope to detect on Earth is bounded by $h_0 \lesssim 4 \times 10^{-24}$. The argument is a modification of an observation due to Blandford (which was unpublished, but credited to him in Thorne's review in [33]).

The argument begins by assuming, very optimistically, that all neutron stars in the Galaxy are born at a very high spin rate and then spin down principally due to gravitational wave emission. For simplicity we shall also assume that all neutron stars follow the same spin-down law $\dot{\nu}(\nu)$

or equivalently $\dot{f}(f)$, although this turns out to be unnecessary to the conclusion. It is helpful to express the spin-down law in terms of the spin-down time scale

$$\tau_{\rm gw}(f) \equiv \frac{f}{|4\dot{f}(f)|}. \tag{8}$$

For a neutron star with constant ellipticity, $\tau_{\rm gw}(f)$ is the time for the gravitational wave frequency to drift down to f from some initial, much higher spin frequency. This time scale is independent of ellipticity and emission mechanism, so long as the emission is quadrupolar. (It is similar to the characteristic age $f/|2\dot{f}|$ used in pulsar astronomy, except that the 2 is replaced by 4 as appropriate for quadrupole rather than dipole radiation.) A source's gravitational wave amplitude h_0 is then related to $\tau_{\rm gw}(f)$ by

$$h_0(f) = d^{-1}\sqrt{\frac{5GI_{zz}}{8c^3\tau_{\rm gw}(f)}}. \tag{9}$$

Here we are assuming that the star is not accreting, so that the angular momentum loss to GWs causes the star to slow down. The case of accreting neutron stars is dealt with separately, below.

We now consider the distribution of neutron stars in space and frequency. Let $N(f)\Delta f$ be the number of Galactic neutron stars in the frequency range $[f - \Delta f/2, f + \Delta f/2]$. We assume that the birthrate has been roughly constant over a long enough time scale that this distribution has settled into a statistical steady state: $dN(f)/dt = 0$ above the minimum frequency $f_{\rm min}$ of our search. (This is not true for millisecond pulsars; see below.) Then $N(f)\dot{f}$ is just the neutron star birthrate $1/\tau_{\rm b}$, where $\tau_{\rm b}$ may be as short as 30 years. For simplicity, we model the spatial distribution of neutron stars in our Galaxy as that of a uniform cylindrical disk, with radius $R_G \approx 10$ kpc and height $H \approx 600$ pc. Then the spatial density $n(f)$ of neutron stars near the Earth, in the frequency range $[f - \Delta f/2, f + \Delta f/2]$, is just $n(f)\Delta f = (\pi R_G^2 H)^{-1} N(f)\Delta f$.

Let $\hat{N}(f, d)$ be that portion of $N(f)$ due to neutron stars whose distance from Earth is less than d. For $H/2 \lesssim d \lesssim R_G$, we have

$$\frac{d\hat{N}(f, d)}{d(d)} = 2\pi d H n(f) \tag{10}$$

$$= 2N(f)\frac{d}{R_G^2} \tag{11}$$

(and it drops off rapidly for $d \gtrsim R_G$). Changing variables from d to h_0 using Eqs. (8) and (9), we have

$$\frac{d\hat{N}(f, h_0)}{dh_0} = \frac{5GI_{zz}}{c^3\tau_b R_G^2} f^{-1} h_0^{-3}. \tag{12}$$

Note that the dependence on the poorly known $\tau_{gw}(f)$ has dropped out of this equation. This was the essence of Blandford's observation.

Now consider a search for GW pulsars in the frequency range $[f_{\rm min}, f_{\rm max}]$. Integrating the distribution in Eq. (12) over this band, we obtain the distribution of sources as a function of h_0:

$$\frac{dN_{\rm band}}{dh_0} = \frac{5GI_{zz}}{c^3\tau_b R_G^2} h_0^{-3} \ln\left(\frac{f_{\rm max}}{f_{\rm min}}\right). \tag{13}$$

The amplitude $h_0^{\rm max}$ of the strongest source is implicitly given by

$$\int_{h_0^{\rm max}}^{\infty} \frac{dN_{\rm band}}{dh_0} dh_0 = \frac{1}{2}. \tag{14}$$

That is, even given our optimistic assumptions about the neutron star population, there is only a 50% chance of seeing a source as strong as $h_0^{\rm max}$. The integral in Eq. (14) is trivial; it yields

$$h_0^{\rm max} = \left[\frac{5GI_{zz}}{c^3\tau_b R_G^2} \ln\left(\frac{f_{\rm max}}{f_{\rm min}}\right)\right]^{1/2}. \tag{15}$$

Inserting $[\ln(f_{\rm max}/f_{\rm min})]^{1/2} \approx 1$ (appropriate for a typical broadband search, as conducted here), and adopting as fiducial values $I_{zz} = 10^{45}$ g cm^2, $R_G = 10$ kpc , and $\tau_b = 30$ yr, we arrive at

$$h_0^{\rm max} \approx 4 \times 10^{-24}. \tag{16}$$

This is what we aimed to show.

We now address the robustness of some assumptions in the argument. First, the assumption of a universal spin-down function $\tau_{\rm gw}(f)$ was unnecessary, since $\tau_{\rm gw}(f)$ disappeared from Eq. (12) and the subsequent equations that led to $h_0^{\rm max}$. Had we divided neutron stars into different classes labeled by i and assigned each a spin-down law $\tau_{\rm gw}^i(f)$ and birthrate $1/\tau_b^i$, each would have contributed its own term to $d\hat{N}/dh_0$ which would have been independent of $\tau_{\rm gw}^i$ and the result for $h_0^{\rm max}$ would have been the same.

Second, in using Eq. (10), we have in effect assumed that the strongest source is in the distance range $H/2 \lesssim d \lesssim R_G$. We cannot evade the upper limit by assuming that the neutron stars have extremely long spin-down times (so that $d < H/2$) or extremely short ones (so that the brightest is outside our Galaxy, $d > R_G$). If the brightest sources are at $d < H/2$ (as happens if these sources have long spin-down times, $\tau_{\rm gw} \gtrsim \tau_b(2R_G/H)^2$), then our estimate of $h_0^{\rm max}$ only decreases, because at short distances the spatial distribution of neutron stars becomes approximately spherically symmetric instead of planar and the right-hand sides of Eqs. (10) and (12) are multiplied by a factor $2r/H < 1$. On the other hand, if $\tau_{\rm gw}(f)$ (in the LIGO range) is much shorter than $\tau_{\rm b}$, then the probability that such an object exists inside our Galaxy is $\ll 1$. For example, a neutron star with $\tau_{\rm gw}(f) = 3$ yr located at $r = 10$ kpc would have $h_0 = 4.14 \times 10^{-24}$, but the probability of currently having

a neutron star with this (or shorter) $\tau_{\rm gw}$ is only $\tau_{\rm gw}/\tau_{\rm b} \lesssim 1/10$.

Third, we have implicitly assumed that each neutron star spins down only once. In fact, it is clear that some stars in binaries are "recycled" to higher spins by accretion, and then spin down again. This effectively increases the neutron star birth rate (since for our purposes the recycled stars are born twice), but since the fraction of stars recycled is very small the increase in the effective birth rate is also small.

2. Expected sensitivity of the S2 search

Typical noise levels of LIGO during the S2 run were approximately $[S_h(f)]^{1/2} \approx 3 \times 10^{-22}$ Hz$^{-1/2}$, where S_h is the strain noise power spectral density, as shown in Fig. 1. Even for a *known* GW pulsar with an average sky position, inclination angle, polarization, and frequency, the amplitude of the signal that we could detect in Gaussian stationary noise with a false alarm rate of 1% and a false dismissal rate of 10% is [19]

$$\langle h_0(f) \rangle = 11.4 \sqrt{\frac{S_h(f)}{T_{\rm obs}}}, \tag{17}$$

where $T_{\rm obs}$ is the integration time and the angled brackets indicate an average source. In all-sky searches for pulsars with *unknown* parameters, the amplitude h_0 must be several times greater than this to rise convincingly above the background. Therefore, in $T_{\rm obs} = 10$ hours of S2 data, signals with amplitude h_0 below about 10^{-22} would not be detectable. This is a factor ≈ 25 greater than the value of $h_0^{\rm max}$ shown in (16), so our S2 analysis is unlikely to be sensitive enough to reveal previously unknown pulsars.

The sensitivity of our search is further restricted by the template bank, which does not include the effects of signal spin-down for reasons of computational cost. Phase mismatch between the signal and matched filter causes the detection statistic (see Sec. VA) to decrease rapidly for GW frequency derivatives $\dot{f}$ that exceed

$$\max[\dot{f}] = \frac{1}{2} T_{\rm obs}^{-2} = 4 \times 10^{-10} \left(\frac{T_{\rm obs}}{10\ {\rm h}} \right)^{-2} \ {\rm Hz\,s^{-1}}. \tag{18}$$

Assuming that all of the spin-down of a neutron star is due to gravitational waves (from a mass quadrupole deformation), our search is restricted to pulsars with ellipticity ϵ less than

$$\epsilon_{\rm sd} = \left(\frac{5c^5 \max[\dot{f}]}{32\pi^4 G I_{zz} f^5} \right)^{1/2}. \tag{19}$$

This limit, derived from combining the quadrupole formula for GW luminosity

$$\frac{dE}{dt} = \frac{1}{10} \frac{G}{c^5} (2\pi f)^6 I_{zz}^2 \epsilon^2 \tag{20}$$

with the kinetic energy of rotation

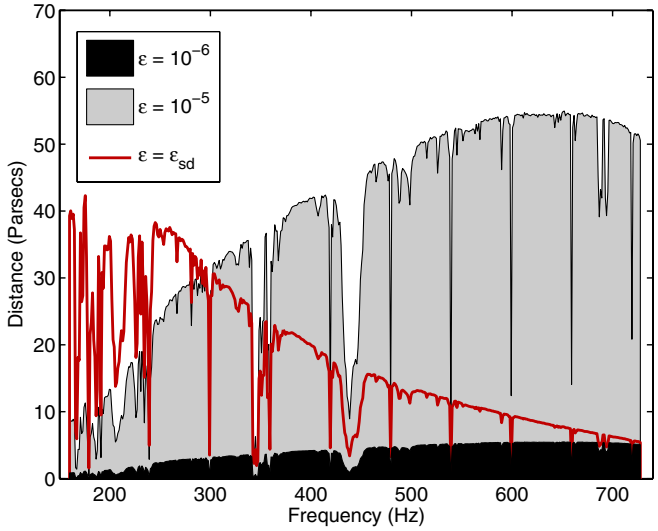

FIG. 2 (color online). Effective average range (defined in the text) of our search as a function of frequency for three ellipticities: 10^{-6} (maximum for a normal neutron star), 10^{-5} (maximum for a more optimistic object), and ϵ_{sd}, the spin-down limit defined in the text. Note that for sources above 300 Hz the reach of the search is limited by the maximum spin-down value of a signal that may be detected without loss of sensitivity.

$$E = \tfrac{1}{2} \pi^2 f^2 I_{zz}, \tag{21}$$

(assuming $f = 2\nu$) takes the numerical value

$$\epsilon_{\rm sd} = 9.6 \times 10^{-6} \left(\frac{10^{45}\ {\rm g\,cm^2}}{I_{zz}} \right)^{1/2} \left(\frac{300\ {\rm Hz}}{f} \right)^{5/2} \tag{22}$$

for our maximum $\dot{f}$.

The curves in Fig. 2 are obtained by combining Eqs. (2) and (17)[1] and solving for the distance d for different values of the ellipticity, using an average value for the noise in the detectors during the S2 run. The curves show the average distance, in the sense of the definition (17), at which a source may be detected.

The black region shows that a GW pulsar with $\epsilon = 10^{-6}$ could be detected by this search only if it were very close, less than ~5 parsecs away. The light gray region shows the distance at which a GW pulsar with $\epsilon = 10^{-5}$ could be detected if templates with sufficiently large spin-down values were searched. However, *this* search can detect such pulsars only below 300 Hz, because above 300 Hz a GW pulsar with $\epsilon = 10^{-5}$ spins down too fast to be

[1]Note that the value of h_0 derived from Eq. (17) yields a value of the detection statistic $2\mathcal{F}$ for an average source as seen with a detector at S2 sensitivity and over an observation time of 10 hours, of about 21, which is extremely close to the value of 20 which is used in this analysis as the threshold for registering candidate events. Thus combining Eqs. (2) and (17) determines the smallest amplitude that our search pipeline could detect (corresponding to a signal just at the threshold), provided appropriate follow-up studies of the registered events ensued.

detected with the no-spin-down templates used. The thick line indicates the distance limit for the (frequency-dependent) maximum value of epsilon that could be detected with the templates used in this search. At certain frequencies below 300 Hz, a GW pulsar could be seen somewhat farther away than 30 pc, but only if it has $\epsilon > 10^{-5}$. Although ϵ_{sd} and the corresponding curve were derived assuming a quadrupolar deformation as the emission mechanism, the results would be similar for other mechanisms. Equation (21) includes an implicit factor $f^2/(2\nu)^2$, which results in ϵ_{sd} and the corresponding range (for a fixed GW frequency f) being multiplied by $f/(2\nu)$, which is $1/2$ for free precession and about $2/3$ for r-modes. Even for a source with optimum inclination angle and polarization, the range increases only by a factor ≈ 2. The distance to the nearest known pulsar in the LIGO frequency band, PSR J0437 − 4715, is about 140 pc. The other nearest neutron stars are at comparable distances [32,34] including RX J1856.5 − 3754, which may be the nearest of all and was found to have a pulsation period out of the LIGO band after this article was submitted [35]. Therefore our search would be sensitive only to previously unknown objects.

While we have argued that a detection would be very unlikely, it should be recalled that Eq. (16) was based on a statistical argument. It is always possible that there is a GW-bright neutron star that is much closer to us than would be expected from a random distribution of supernovae (for example due to recent star formation in the Gould belt as considered in [36]). It is also possible that a "blind" search of the sort performed here could 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. With respect to S2, this is a factor 10 improvement in the amplitude strain noise at most frequencies. The greater length of the data set will also increase the sensitivity to pulsars by a factor of a few (the precise value depends on the combination of coherent and incoherent analysis methods used). The net result is that initial LIGO will have h_0 reduced from the S2 value by a factor of 30 or more to a value comparable to $h_0^{\max} \approx 4 \times 10^{-24}$ of Eq. (16).

C. Accreting neutron stars

1. Maximum expected signal amplitude at earth

The robust upper limit in Eq. (16) refers only to non-accreting neutron stars, since energy conservation plays a crucial role. If accretion replenishes the star's angular momentum, a different but equally robust argument (i.e., practically independent of the details of the emission mechanism) can be made regarding the maximum strain $h_0^{\max}$ at the Earth. In this case $h_0^{\max}$ is set by the x-ray luminosity of the brightest x-ray source.

The basic idea is that if the energy (or angular momentum) lost to GWs is replenished by accretion, then the strongest GW emitters are those accreting at the highest rate, near the Eddington limit. Such systems exist: the low-mass x-ray binaries (LMXBs), so-called since the accreted material is tidally stripped from a low-mass companion star. The accreted gas hitting the surface of the neutron star is heated to 10^8 K and emits x-rays. As noted several times over the years [1,37,38], if one assumes that spin-down from GW emission is in equilibrium with accretion torque, then the GW amplitude h_0 is directly related to the x-ray luminosity:

$$h_0 \approx 5 \times 10^{-27} \left(\frac{300\ \text{Hz}}{\nu}\right)^{1/2} \left(\frac{F_x}{10^{-8}\ \text{erg cm}^{-2}\,\text{s}^{-1}}\right)^{1/2}, \tag{23}$$

where F_x is the x-ray flux. In the 1970s when this connection was first proposed, there was no observational support for the idea that the LMXBs are strong GW emitters. But the spin frequencies of many LMXBs are now known, and most are observed to cluster in a fairly narrow range of spin frequencies 270 Hz $\lesssim \nu \lesssim$ 620 Hz [39]. Since most neutron stars will have accreted enough matter to spin them up to near their theoretical maximum spin frequencies, estimated at ~1400 Hz, the observed spin distribution is hard to explain without some competing mechanism, such as gravitational radiation, to halt the spin-up. Since the gravitational torque scales as ν^5, gravitational radiation is also a natural explanation for why the spin frequencies occupy a rather narrow window: a factor 32 difference in accretion rate leads to only a factor 2 difference in equilibrium spin rate [1].

If the above argument holds, then the accreting neutron star brightest in x-rays is also the brightest in gravitational waves. Sco X-1, which was the first extrasolar x-ray source discovered, is the strongest persistent x-ray source in the sky. Assuming equilibrium between GWs and accretion, the gravitational wave strain of Sco X-1 at the Earth is

$$h_0 \approx 3 \times 10^{-26} \left(\frac{540\ \text{Hz}}{f}\right)^{1/2}, \tag{24}$$

which should be detectable by second generation interferometers. The gravitational wave strains from other accreting neutron stars are expected to be lower.

2. Expected sensitivity of S2 search for Sco X-1

The orbital parameters of Sco X-1 are poorly constrained by present (mainly optical) observations and large uncertainties affect the determination of the rotation frequency of the source (details are provided in Sec. IV B 2). The immediate implication for a coherent search for gravitational waves from such a neutron star is that a very large number of discrete templates are required to cover the

relevant parameter space, which in turn dramatically increases the computational costs [40]. The optimal sensitivity that can be achieved with a coherent search is therefore set primarily by the length of the data set that one can afford to process (with fixed computational resources) and the spectral density of the detector noise. As we discuss in Sec. IV B 2, the maximum span of the observation time set by the computational burden of the Sco X-1 pipeline (approximately one week on ≈ 100 CPUs) limits the observation span to 6 hours.

The overall sensitivity of the search that we are describing is determined by each stage of the pipeline, which we describe in detail in Sec. V B. Assuming that the noise in the instrument can be described as a Gaussian and stationary process (an assumption which, however, breaks down in some frequency regions and/or for portions of the observation time), we can statistically model the effects of each step of the analysis and estimate the sensitivity of the search. The results of such modelling through the use of Monte Carlo simulations are shown in Fig. 3 where we give the expected upper limit sensitivity of the search. We contrast this with the hypothetical case in which the Sco X-1 parameters are known perfectly making it a single filter target for the whole duration of the S2 run. The dramatic difference (of at least an order of magnitude) between the estimated sensitivity curves of these two scenarios is primarily due to the large parameter space we have to search. This has two consequences, which contribute to degrading the sensitivity of the analysis: (i) we are computationally limited by the vast number of templates that we must search and therefore must reduce the observation to a subsection of the S2 data, and (ii) sampling a large number of independent locations increases the probability that noise alone will produce a high value of the detection statistic.

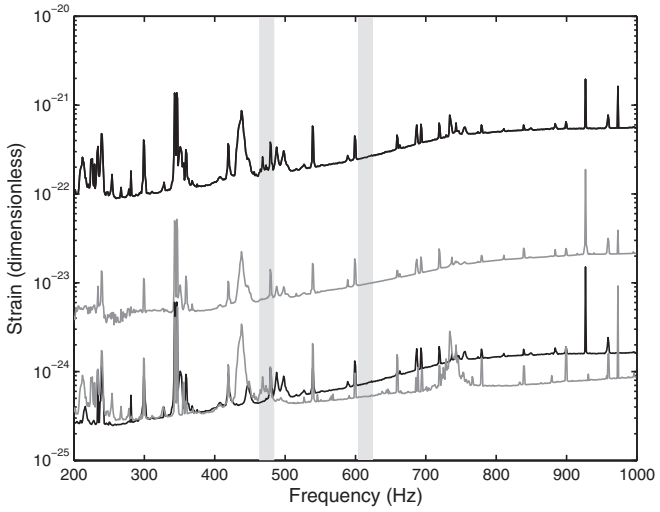

FIG. 3. Here we show the expected upper limit sensitivity of the S2 Sco X-1 search. The upper black curve represents the expected sensitivity of the S2 analysis based on an optimally selected 6-hour data set (chosen specifically for our search band). The gray curve (second from the top) shows the sensitivity in the hypothetical case in which *all* of the Sco X-1 system parameters are known exactly making Sco X-1 a single filter target and the entire S2 data set is analyzed. Both curves are based on a 95% confidence upper limit. The remaining curves represent $\sqrt{S_h(f)/T_{\rm obs}}$ for L1 (black) and H1 (gray); $S_{\rm h}(f)$ is the typical noise spectral density that characterizes the L1 and H1 data, and $T_{\rm obs}$ is the actual observation time (taking into account the duty cycle, which is different for L1 and H1) for each instrument.

We note that the S2 Sco X-1 analysis is a factor of ≈ 5000 less sensitive than the characteristic amplitude given in Eq. (24). In the hypothetical case in which Sco X-1 is a single filter target and we are able to analyze the entirety of S2 data, then we are still a factor ~ 100 away. However, as mentioned in the introduction, the search reported in this paper will be one of the stages of a more sensitive ''hierarchical pipeline'' that will allow us to achieve quasioptimal sensitivity with fixed computational resources.

IV. SIGNAL MODEL

A. The signal at the detector

We consider a rotating neutron star with equatorial coordinates α (right ascension) and δ (declination). Gravitational waves propagate in the direction $\hat{k}$ and the star spins around an axis whose direction, assumed to be constant, is identified by the unit vector $\hat{s}$.

The strain $h(t)$ recorded at the interferometer output at detector time t is

$$h(t) = h_0[\tfrac{1}{2}(1+\cos^2\iota)F_+(t;\alpha,\delta,\psi)\cos\Phi(t) + \cos\iota F_\times(t;\alpha,\delta,\psi)\sin\Phi(t)], \quad (25)$$

where ψ is the polarization angle, defined as $\tan\psi = [\hat{k}\cdot(\hat{s}\times\hat{z})]/[(\hat{s}\cdot\hat{k})(\hat{z}\cdot\hat{k}) - (\hat{s}\cdot\hat{z})]$, $\hat{z}$ is the direction to the north celestial pole, and $\cos\iota = \hat{k}\cdot\hat{s}$. Gravitational wave laser interferometers are all-sky monitors with a response that depends on the source location in the sky and the wave polarization: this is encoded in the (time-dependent) antenna beam patterns $F_{+,\times}(t;\alpha,\delta,\psi)$. The term $\Phi(t)$ in Eq. (25) represents the phase of the received gravitational signal.

The analysis challenge to detect weak quasiperiodic continuous gravitational waves stems from the Doppler shift of the gravitational phase $\Phi(t)$ due to the relative motion between the detector and the source. It is convenient to introduce the following times: t, the time measured at the detector; T, the solar-system-barycenter (SSB) coordinate time; and t_p, the proper time in the rest frame of the pulsar.[2]

[2]Notice that our notation for the three different times is different from the established conventions adopted in the radio pulsar community, e.g. [41].

The timing model that links the detector time t to the coordinate time T at the SSB is

$$T = t + \frac{\vec{r}\cdot\hat{n}}{c} + \Delta_{E\odot} - \Delta_{S\odot}, \tag{26}$$

where $\vec{r}$ is a (time-dependent) vector from the SSB to the detector at the time of the observations, $\hat{n}$ is a unit vector towards the pulsar (it identifies the source position in the sky) and $\Delta_{E\odot}$ and $\Delta_{S\odot}$ are the solar system Einstein and Shapiro time delays, respectively [41]. For an isolated neutron star t_p and T are equivalent up to an additive constant. If the source is in a binary system, as it is the case for Sco X-1, significant additional accelerations are involved, and a further transformation is required to relate the proper time t_p to the detector time t. Following [41], we have

$$T - T_0 = t_p + \Delta_R + \Delta_E + \Delta_S, \tag{27}$$

where Δ_R, the Roemer time delay, is analogous to the solar system term $(\vec{r}\cdot\hat{n})/c$; Δ_E and Δ_S are the orbital Einstein and Shapiro time delay, analogous to $\Delta_{E\odot}$ and $\Delta_{S\odot}$; and T_0 is an arbitrary (constant) reference epoch. For the case of Sco X-1, we consider a circular orbit for the analysis (cf. Sec. IV B 2 for more details) and therefore set $\Delta_E = 0$. Furthermore, the binary is nonrelativistic and from the source parameters we estimate $\Delta_S < 3\ \mu$s which is negligible. For a circular orbit, the Roemer time delay is simply given by

$$\Delta_R = \frac{a_{\mathrm{p}}}{c}\sin(u + \omega), \tag{28}$$

where a_{p} is the radius of the neutron star orbit projected on the line of sight, ω the argument of the periapsis and u the so-called eccentric anomaly; for the case of a circular orbit $u = 2\pi(t_p - t_{p,0})/P$, where P is the period of the binary and $t_{p,0}$ is a constant reference time, conventionally referred to as the "time of periapse passage."

In this paper we consider gravitational waves whose *intrinsic* frequency drift is negligible over the integration time of the searches (details are provided in the next section), both for the blind analysis of unknown isolated neutron stars and Sco X-1. The phase model is simplest in this case and given by

$$\Phi(t_p) = 2\pi f_0 t_p + \Phi_0, \tag{29}$$

where Φ_0 is an overall constant phase term and f_0 is the frequency of the gravitational wave at the reference time.

B. Parameter space of the search

Both searches require exploring a three-dimensional parameter space, consisting of two "position parameters" and the unknown frequency of the signal. For the all-sky blind analysis aimed at unknown isolated neutron stars one needs to Doppler correct the phase of the signal for any given point in the sky, based on the angular resolution of the instrument over the observation time, and so a search is performed on the sky coordinates α and δ. For the Sco X-1 analysis, the sky location of the system is known; however, the system is in a binary orbit with poorly measured orbital elements; thus, one needs to search over a range of orbital parameter values. The frequency search parameter is for both searches the f_0 defined by Eq. (29), where the reference time has been chosen to be the time-stamp of the first sample of each data set. The frequency band over which the two analyses are carried out is also different, and the choice is determined by astrophysical and practical reasons. As explained in Sec. V C, the data set in H1 does not coincide in time with the L1 data set for either of the analyses. Consequently a signal with a nonzero frequency derivative would appear at a different frequency template in each data set. However, for the maximum spin-down rates considered in this search, and given the time lag between the two data sets, the maximum difference between the search frequencies happens for the isolated objects search and amounts to 0.5 mHz. We will see that the frequency coincidence window is much larger than this and that when we discuss spectral features in the noise of the data and locate them based on template-triggers at a frequency f_0, the spectral resolution is never finer than 0.5 mHz. So for the practical purposes of the present discussion we can neglect this difference and will often refer to f_0 generically as the signal's frequency.

1. Isolated neutron stars

The analysis for isolated neutron stars covers the entire sky and the frequency range 160–728.8 Hz. The low frequency end of the band was chosen because the depth of our search degrades significantly below 160 Hz, see Fig. 2. The choice of the high frequency limit at 728.8 is primarily determined by the computational burden of the analysis, which scales as the square of the maximum frequency that is searched for.

In order to keep the computational costs at a reasonable level (<1 month on $\lesssim 800$ CPUs), no explicit search over spin-down parameters was carried out. The length of the data set that is analyzed is approximately 10 hours, thus no loss of sensitivity is incurred for sources with spin-down rates smaller than 4×10^{-10} Hz s^{-1}; see Eq. (18). This is a fairly high spin-down rate compared to those measured in isolated radio pulsars; however, it does constrain the sensitivity for sources above 300 Hz, as can be seen from Fig. 2.

2. Sco X-1

Sco X-1 is a neutron star in a 18.9 h orbit around a low-mass ($\sim 0.42 M_\odot$) companion at a distance $d = 2.8 \pm 0.3$ kpc from Earth. Table I contains a summary of the parameters and the associated uncertainties that are relevant for gravitational wave observations. In this section we summarize the area in parameter space over which the

TABLE I. The parameters of the low-mass x-ray binary Scorpius X-1. The quoted measurement errors are all $1-\sigma$. We refer the reader to the text for details and references.

right ascension	α	16 h 19 m 55.0850 s
declination	δ	$-15°38'24.9''$
proper motion (east-west direction)	μ_x	$-0.006\,88 \pm 0.000\,07$ arcsec yr^{-1}
proper motion (north-south direction)	μ_y	$0.012\,02 \pm 0.000\,16$ arcsec yr^{-1}
distance	d	2.8 ± 0.3 kpc
orbital period	P	$68\,023.84 \pm 0.08$ sec
time of periapse passage	$\bar{T}$	$731\,163\,327 \pm 299$ sec
projected semimajor axis	a_{p}	1.44 ± 0.18 sec
eccentricity	e	$<3 \times 10^{-3}$
QPOs frequency separation	237	± 5 Hz $\leq \Delta\nu_{\mathrm{QPO}} \leq 307 \pm 5$ Hz

analysis is carried out. More details are given in Appendix A.

We will assume the observation time to be 6 hours. This is approximately what was adopted for the analysis presented in this paper and we shall justify this choice at the end of the section.

The position of Sco X-1 (i.e. the barycenter of the binary system) is known to high accuracy and we "point" (in software) at that region of the sky. Of the three parameters that describe the circular orbit of a star in a binary system, the orbital period (P), the projection of the semimajor axis of the orbit a_{p} (which for $e = 0$ corresponds to the projected radius of the orbit), and the location of the star on the orbit at some given reference time, which we define as the *orbital phase reference time* $\bar{T}$, P can be regarded as known over 6-hour integration time and the search therefore requires a discrete grid of filters in the $(a_{\mathrm{p}}, \bar{T})$ space. Notice that we assume that Sco X-1 is in a circular orbit and analyze the data under this assumption. A zero eccentricity is what one expects for a semidetached binary system and is consistent with the best fits of the orbital parameters [42]. Limitations of this assumption are discussed in Appendix A. However, we quantify (for a smaller set of the parameter space) the efficiency of the pipeline in searching for gravitational waves emitted by a binary with non zero eccentricity; in other words, we quote upper limits for different values of the eccentricity that are obtained with nonoptimal search templates.

For the frequency of the gravitational radiation, f_0, we confine the analysis to the two 20 Hz wide bands (464–484 Hz and 604–624 Hz) that bound the range of the drift of ν, according to currently acceptable models for the kHz QPOs.

The total computational time for the analysis can be split into two parts: (i) the *search time* T_{search} needed to search the data and, if no signal is detected, (ii) the *upper limit time* T_{inj} required to repeatedly inject and search for artificially generated signals for the purposes of setting the upper limits. Let T_{span} be the span of the data set which is analyzed, that is the difference between the time stamps of the first and last data point in the time series. Let T_{obs} be the effective duration of the data set containing nonzero data points. The definitions imply $T_{\mathrm{obs}} \leq T_{\mathrm{span}}$, and for data with no gaps $T_{\mathrm{obs}} = T_{\mathrm{span}}$. For a search confined to a period (sufficiently) shorter than the orbital period of the source, the two computational times are

$$T_{\mathrm{search}} \approx 90\ \mathrm{hrs} \times \left(\frac{\Delta f}{40\ \mathrm{Hz}}\right)\left(\frac{\Delta\bar{T}}{598\ \mathrm{s}}\right)\left(\frac{0.1}{\mu}\right)^{3/2}\left(\frac{100}{N_{\mathrm{cpu}}}\right) \times \frac{1}{2}\sum_{L1,H1}\left(\frac{T_{\mathrm{span}}}{6\ \mathrm{hrs}}\right)^{7}\left(\frac{T_{\mathrm{obs}}}{T_{\mathrm{span}}}\right), \tag{30}$$

$$T_{\mathrm{inj}} \approx 55\ \mathrm{hrs} \times \left(\frac{N_{\mathrm{trials}}}{5000}\right)\left(\frac{N_{h_0}}{20}\right)\left(\frac{100}{N_{\mathrm{cpu}}}\right)\frac{1}{2}\sum_{L1,H1}\left(\frac{T_{\mathrm{span}}}{6\ \mathrm{hrs}}\right)\left(\frac{T_{\mathrm{obs}}}{T_{\mathrm{span}}}\right), \tag{31}$$

where Δf is the search frequency band, $\Delta\bar{T}$ the search range for the time of periapse passage, μ is the template bank mismatch, and N_{cpu} is the number of $\sim$2 GHz CPUs available [43]. The quantities N_{trials} and N_{h_0} are the number of artificial signals injected per value of h_0 and the number of different values of h_0 injected, respectively. Note the steep dependency of the search time T_{search} on the maximum observation time span T_{span}. The contributing factors to this scaling are the increasing number of orbital and frequency filters, N_{orb} and N_{freq}, respectively, with observation time span, where $N_{\mathrm{orb}} \propto T_{\mathrm{span}}^5$ and $N_{\mathrm{freq}} \propto T_{\mathrm{obs}}$. There is also a linear scaling of computational time with T_{span} (corrected by the factor $T_{\mathrm{obs}}/T_{\mathrm{span}}$ that takes into account only the nonzero data points) due to increased data volume being analyzed. From Eqs. (30) and (31) it is therefore clear that if one wants to complete the full analysis over a period $\lesssim 1$ week the choice $T_{\mathrm{span}} = 6$ h is appropriate.

V. ANALYSIS OF THE DATA

The inner core of the analysis is built on the frequency-domain matched-filter approach that we applied to the data collected during the first science run to place an upper limit on gravitational radiation from PSR J1939 + 2134 [19].

However, this analysis is considerably more complex with respect to [19] because (i) the search is carried out over a large number of templates (either over sky position or source orbital parameters), (ii) coincidences are looked for in the output of the searches on two interferometers in order to reduce the false alarm probability and thereby improve the overall sensitivity of the search, and (iii) the upper limit is derived from the maximum *joint* significance of coincident templates.

A. The detection statistic

The optimal detection statistic (in the maximum likelihood sense) to search coherently for quasimonochromatic signals is the so-called $\mathcal{F}$-statistic[3] introduced in [44]. This statistic can be extended in a straightforward manner to the case of a signal from a pulsar in a binary system.

In the absence of signal, $2\mathcal{F}$ is distributed according to a (central) χ^2 distribution with 4 degrees of freedom and the relevant probability density function is given by

$$p_0(2\mathcal{F}) = \frac{2\mathcal{F}}{4} e^{-(2\mathcal{F}/2)}. \tag{32}$$

We define the false alarm probability of $2\mathcal{F}$ as

$$P_0(2\mathcal{F}) = \int_{2\mathcal{F}}^{\infty} p_0(2\mathcal{F}')d(2\mathcal{F}'). \tag{33}$$

In the presence of a signal, $2\mathcal{F}$ follows a noncentral χ^2 distribution with 4 degrees of freedom and noncentrality parameter ρ^2; the associated probability density function is

$$p_1(2\mathcal{F}) = \frac{1}{2} e^{-(2\mathcal{F}+\rho^2)/2} \sqrt{\frac{2\mathcal{F}}{\rho^2}} I_1(\sqrt{2\mathcal{F}\rho^2}), \tag{34}$$

where I_1 is the modified Bessel function of the first kind of order one and

$$\rho^2 = \frac{2}{S_h(f)} \int_0^{T_{\rm obs}} h^2(t)dt. \tag{35}$$

The expected value of $2\mathcal{F}$ is $4 + \rho^2$. From Eq. (35) it is clear that the detection statistic is proportional to the square of the amplitude of the gravitational wave signal, h_0^2, given by Eq. (2).

B. Pipeline

The search pipeline is schematically illustrated in Fig. 4. A template bank is set up for each search covering the parameter space under inspection. For both analyses the template bank is three dimensional: it covers right ascension and declination for the unknown isolated pulsar search, and the orbital phase reference time and the projection of the orbital semimajor axis for the Sco X-1 analysis. In addition, in both cases we search for the unknown gravitational wave frequency.

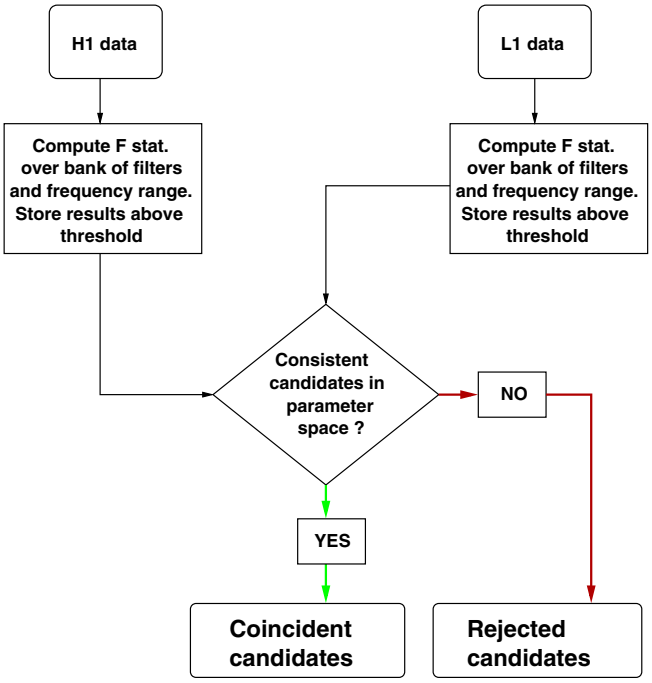

FIG. 4 (color online). Workflow of the pipeline.

The data stream is treated in exactly the same way for each search: the full search frequency band is divided into smaller (~ 1 Hz) *sub-bands*,[4] the $\mathcal{F}$-statistic is computed at every point in the template bank, and lists of candidate templates are produced. Search template values are recorded when the detection statistic exceeds the value $2\mathcal{F} = 20$, and we will refer to them as *registered templates*. Note that we will also refer to these templates as "events," by analogy with the time-domain matched filtering analysis.

In the search conducted in [19], we searched a single template in four detectors. Here we search a total of 5×10^{12} templates in each detector for the isolated pulsars and 3×10^{10} templates overall in two detectors for Sco X-1. In order to reduce the number of *recorded* templates, only the maximum of the detection statistic over a frequency interval at fixed values of the remaining template parameters is stored. The frequency interval over which this maximization is performed is based on the maximum expected width of the detection statistic for an actual signal.

[3]We would like to stress that this statistic is completely unrelated to the F-statistic described in statistical textbooks to test the null hypothesis for two variances drawn from distributions with the same mean.

[4]Each of the sub-bands corresponds to the frequency region which the loudest candidate over the sky or the orbital parameters is maximized over and should be of comparable width with respect to typical noise floor variations. The precise bandwidth size is dictated by convenience in the computational setup of the analyses, reflecting a different distribution of the computational load among the various nodes of the computer clusters used.

For each recorded template from one detector, the list of recorded templates from the other detector is scanned for template(s) which are close enough so that an astrophysical signal could have been detected in both templates. The criteria used to define ''closeness'' are different for the two searches and will be described in Secs. V F 1 and V F 2. This procedure yields a third list of templates that are what we refer to as the *coincident templates*. These are template values for which the $\mathcal{F}$-statistic is above threshold *and* such that they could be ascribed to the same physical signal in both data streams. The coincident templates are then ranked according to their *joint significance*.

The *most significant coincident template* is identified in each $\sim$1 Hz sub-band; we also refer to this as the *loudest event* for that frequency sub-band. An upper limit on the value of h_0 from a population of isolated sources or of Sco X-1 systems is placed in each frequency sub-band based on its loudest coincident event. Following [19], this is done by injecting in the real data a set of fake signals at the same level of significance as the loudest measured event and by searching the data with the same pipeline as was used in the analysis. The upper limit procedure is described in Sec. V G.

C. Selection of the data set

The data input to the search is in the form of short time baseline Fourier transforms (SFTs) of chunks of the data. At fixed observation time the computational cost of a search increases linearly with the number of short Fourier transforms employed. Hence, the longer the time baseline of the SFT, the less computationally intensive the search. There are two constraints to making the SFT time baseline long: (i) the noise is estimated on the SFT time scale and thus it should be reasonably stationary on such time scale; and (ii) the signal-to-noise ratio of a putative source will be significantly degraded if the Doppler modulation during the SFT time baseline is of the order of $1/T_{\mathrm{SFT}}$. For the S2 data set and a search extending to about 750 Hz we chose 30 minutes as the time baseline for the SFTs of the search for signals from isolated sources. The Sco X-1 search was carried out using 60 s long SFTs due to the more significant acceleration produced by the orbital motion.

As described in [19] the SFT data is normalized by the noise spectral amplitude. This quantity is estimated for each SFT from the actual data near to the frequency bin of interest. In [19] we used a simple average over frequencies around the target search frequency. That approach worked well because in the vicinity of the target frequency there were no spectral disturbances. But clearly we cannot count on this being the case while searching over several hundred Hz. We have thus adopted a spectral running median estimate method [45–47]. In the isolated pulsar search we have chosen a very conservative window size of 50 1800 s-baseline-SFT bins (27.8 mHz) corresponding to a little under twice the number of terms used in the demodulation routine that computes the detection statistic through the integrals (108) and (109) of [44]. We estimate the noise at every bin as the median computed on $25 + 25 + 1$ values, corresponding to the 25 preceding bins, the bin itself, and the 25 following bins. If an outlier in the data were due to a signal, our spectral estimate would be insensitive to it, and thus we would be preserving it in the normalized data. A window size of 50 60 s-baseline-SFT frequency bins (0.833 Hz) was also used for the upper frequency band of the Sco X-1 search, 604–624 Hz. Because of the presence of some large spectral features in the lower band, 464–484 Hz, a window size of 25 60 s-baseline-SFT bins (0.417 Hz) was used in an attempt to better track the noise floor. Noise disturbances are evident in Fig. 7, where we show (with frequency resolution $1/60$ Hz) the average noise spectral density of the data set used in the analysis. Notice that the lower frequency band presents numerous spectral features, especially in H1; moreover, a strong and broad (approximately 2 Hz) excess noise in both detectors is evident around 480 Hz, which corresponds to a harmonic of the 60 Hz powerline.

The reconstruction of the strain from the output of the interferometer is referred to as the calibration. Details regarding the calibration for the S2 run can be found in [27]. Both analyses presented here use a calibration performed in the frequency domain on SFTs of the detector output. The SFT-strain $h(f)$ is computed by constructing a response function $R(f, t)$ that acts on the interferometer output $q(f)$: $h(f) = R(f, t)q(f)$. Due mainly to changes in the amount of light stored in the Fabry-Perot cavities of the interferometers, the response function, $R(f, t)$, varies in time. These variations are measured using sinusoidal excitations injected into the instrument. Throughout S2, changes in the response were computed every 60 seconds. The SFTs used were 30 minutes long for the isolated pulsar analysis and an averaging procedure was used to estimate the response function for each SFT. For the binary search, which uses 60 s SFTs, a linear interpolation was used, since the start times of the SFTs do not necessarily correspond to those at which the changes in the response were measured.

The observation time chosen for the two searches is significantly less than the total observation time of S2, due to computational cost constraints: about 10 hours and 6 hours for the isolated pulsar and Sco X-1 searches, respectively. We picked the most sensitive data stretches covering the chosen observation times; the criteria used to select the data sets are described below, and the differences reflect the different nature of the searches.

1. Data selection for the isolated neutron star search

Since the blind search for isolated neutron stars is an all-sky search, the most sensitive data are chosen based only on the noise performance of the detectors. The sensitivity is evaluated as an average of the sensitivity at different

frequencies in the highest sensitivity band of the instrument. In particular, the noise is computed in six sub-bands that span the lowest 300 Hz range to be analyzed. The sub-bands are 1 Hz wide, with lowest frequencies, respectively, at 162 Hz, 219 Hz, 282 Hz, 338 Hz, 398 Hz, and 470 Hz. These sub-bands were chosen in regions free of spectral disturbances and the average power in these frequency regions can be taken as a measure of the noise floor there. Even though the search band extends up to 730 Hz, we have chosen these reference sub-bands to lie below 500 Hz, because this is the most sensitive frequency range of our instruments. We construct sets of 20 SFTs (10 hours of data), with the constraint that the data employed in each set does not span more than 13 hours. This constraint stems from computational requirements: the spacing used for the template grid in the sky shrinks very fast with increasing *spanned* observation time. If the data contains no gaps then each 10-hour set differs from the previous only by a single SFT. For H1 we are able to construct 892 such sets, for L1 only 8, reflecting the rather different duty factor in the two instruments. This is obvious from the plots of Fig. 5: For H1 we were able to cover with sets of 20 SFTs the entire run in a fairly uniform way. For L1 it was possible to find sets of nearly contiguous SFTs only in the first and second quarter of the run. We finally compute the average over the different frequency sub-bands and we pick the set for which this number is the smallest. Figure 5 shows this average over the frequency bands and the cross points to the lowest-noise SFT set.

The data sets chosen were for H1 20 30-minute SFTs starting at GPS time 733 803 157 that span 10 hours, and

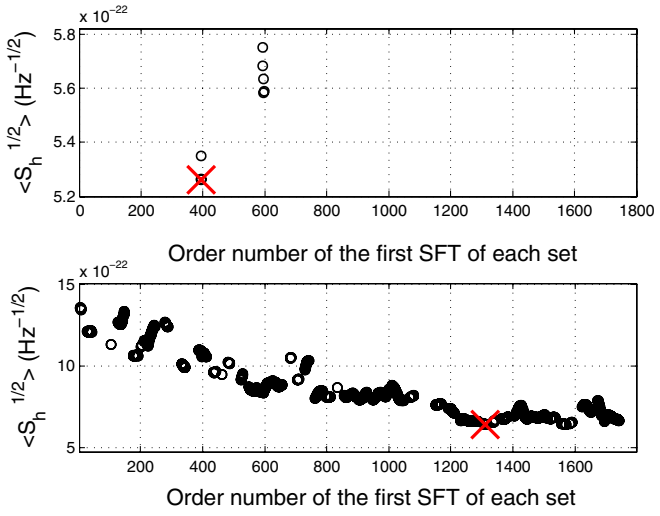

FIG. 5 (color online). The average of the noise over various 1 Hz sub-bands as described in the text for different sets of 20 SFTs from data of the L1 detector (top plot) and H1 detector (bottom plot). The x-axis labels the order number of the first SFT in each set. SFT #1 is the first SFT of the run. Neighboring sets only differ by one SFT. The cross indicates the data set chosen for the search for signals from isolated objects.

for L1 20 30-minute SFTs starting at GPS time 732 489 168 spanning 12.75 hours. The plots of Fig. 6 show the average power spectral density of this data set for the two detectors separately (top two plots) and the average of the same data over 1.2 Hz wide sub-bands and over the two detectors (bottom plot).

2. Data selection for the Sco X-1 search

We choose to analyze in each detector the most sensitive S2 data set which does not span more than 6 hours, which we have set based on computational cost constraints. To rank the sensitivity of a data set we use the figure of merit

$$Q = \sqrt{\frac{5\langle S_h\rangle}{[A+B]T_{\rm obs}}}. \tag{36}$$

Note that Q, S_h, A, and B depend on the time $T_{\rm start}$ corresponding to the first data point of a given data segment that we may wish to analyze, on the total span of the observation (including data gaps) $T_{\rm span}$ and on the effective time containing nonzero data points $T_{\rm obs}$. $\langle S_h\rangle$ is the noise spectral density averaged over the frequency search bands and the data set. The two quantities A and B are the integrals of the amplitude modulation factors and take into account the change of sensitivity of the instruments for the Sco X-1 location in the sky as a function of the time at which the observation takes place (explicit expressions for A and B are given in [44]). In our calculation of Q we take into account the presence of data gaps over $T_{\rm span}$, and we average over the unknown angles ι and ψ. From Eq. (35) it is straightforward to recognize that Q^2 is simply related to the noncentrality parameter ρ^2 for a signal amplitude h_0 by

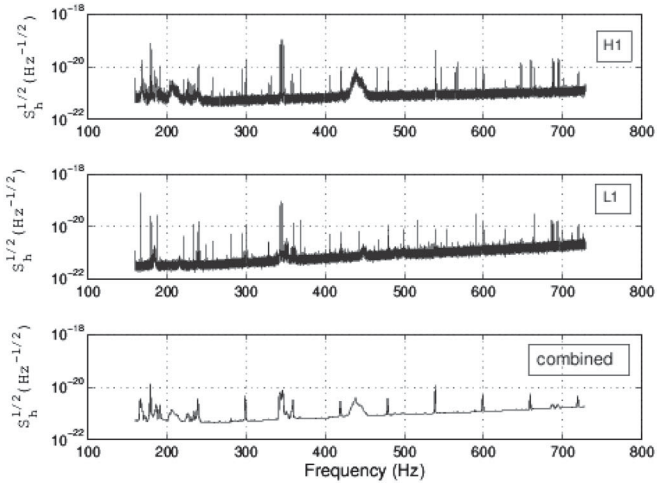

FIG. 6. The average amplitude spectral density $\sqrt{S_{\rm h}(f)}$ of the data of the two detectors used for the isolated pulsars analysis. The bottom plot is the average over 1.2 Hz wide sub-bands of the average of the top two plots. The frequency resolution of the top two plots is $\frac{1}{1800}$ Hz. The frequency resolution of the bottom plots is 2160 times coarser.

$$\langle \rho^2 \rangle_{\iota,\psi} = \frac{[A + B]T_{\mathrm{obs}}}{5\langle S_h \rangle} h_0^2. \tag{37}$$

The parameter Q is therefore a faithful measure of the sensitivity of a given data set for the Sco X-1 search: it combines the effects of the variation with time of the detectors' noise level, duty cycle, and the (angle averaged) sensitivity to the specified sky position. Note that for $T_{\mathrm{span}} \lesssim 1$ day, the location of the source is a strong factor in the choice of the optimal data set; for longer observation times the quantity $A + B$ becomes constant. By tuning the choice of the data set to exactly the Sco X-1 sky position, we have achieved a gain in sensitivity $\gtrsim 2$ compared to selecting the data set only based on the noise level.

We compute the figure of merit Q for all possible choices of data segments with $T_{\mathrm{span}} \leq 6$ hours; the values of Q for the whole S2 are shown in Fig. 8. The data sets that we select for the analysis span 21 611 s with 359 SFTs for H1, and 190 SFTs spanning 18 755 s for L1. Notice that T_{span} is different for the two detectors (which has an impact on the choice of the orbital template banks for the two instruments), and that the L1 and H1 data sets are not coincident in time; based on the relatively short observation time and therefore coarse frequency resolution of the search the maximum spin-up/spin-down of the source due to accretion would change the signal frequency by only ~ 0.1 frequency bins, which is negligible.

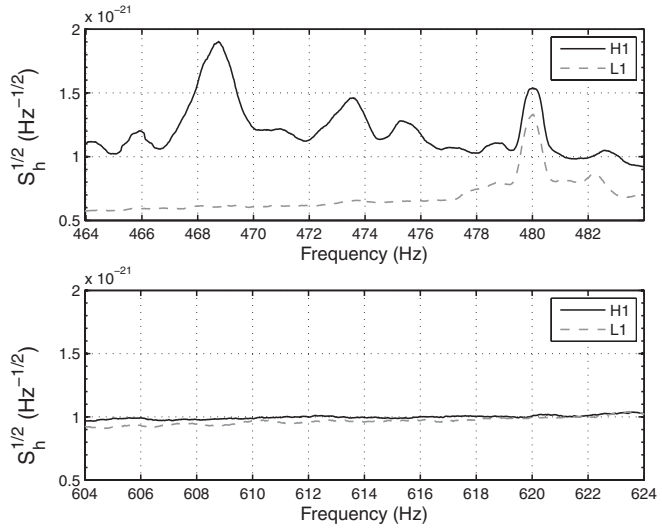

FIG. 7. The amplitude spectral density $\sqrt{S_{\mathrm{h}}(f)}$ of the optimally chosen data sets for L1 and H1 and for both frequency bands averaged over each 60 sec SFT. The solid black and dashed gray lines correspond to H1 and L1, respectively. Note that the lower band contains a feature common to both detectors, a 60 Hz powerline harmonic at 480 Hz with a width of ~ 2 Hz. The H1 data set also contains a variety of other features in the lower band, some equally as large as the powerline harmonic. The upper band is comparatively clean with no visible features.

D. Template banks

In this section we describe the construction of the template (or filter) banks used in the searches. More details can be found in Appendix B. The optimal strategy for laying a filter bank is through a metric approach (cf. [10,48]), and this is used for the orbital parameter grid employed in the Sco X-1 analysis. By contrast, the search for signals from isolated objects uses a suboptimal grid (in terms of computational efficiency, but not for the purpose of recovering signal-to-noise ratio) to cover the sky; a full metric approach was not developed for this search at the time that this analysis was performed.

1. Isolated neutron stars

Two independent grids are employed: one for sky position and one for frequency. The grid in frequency is uniform with a spacing $\Delta f_0 = 3.472 \times 10^{-6}$ Hz which is about a factor of 8 smaller than the inverse of the observation time. To cover the sky we choose an isotropic grid with equatorial spacing of 0.02 rad. Such a grid covers the celestial sphere with just under 31 500 patches of approximately equal surface area. The number of templates in right ascension α at any given declination δ is proportional to $\cos\delta$. At fixed α the spacing in δ is constant, and equal

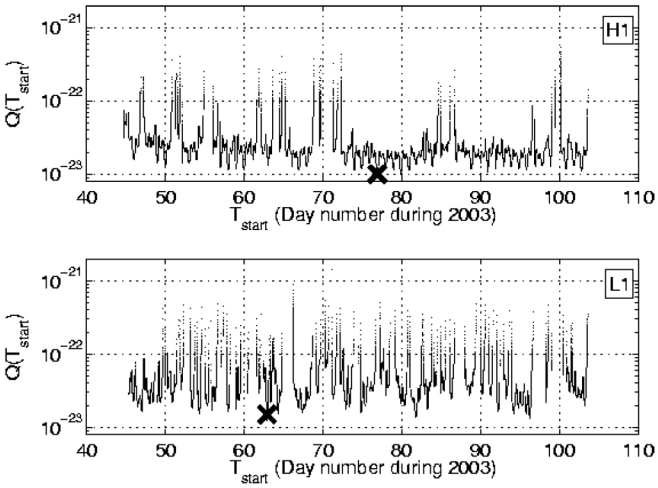

FIG. 8. The sensitivity of the LIGO interferometers during S2 for a search targeted for Sco X-1. The plots show the evolution of the sensitivity quality factor Q, Eq. (36), as a function of observation start time; each point corresponds to a maximum observation time span of 6 hours. Because of intermittent loss of lock during the S2 run each 6-hour span can contain significantly less than 6 hours of data. Note that H1's sensitivity appears more consistent than that of L1. This is due to the lower variation in the H1 noise during the run and the 74% duty cycle compared to L1's 37%. The periodic structures, more visible in the H1 curve, are caused by the daily variation of the detectors' antenna pattern due to the Earth's spin. The crosses indicate the start times corresponding to the best sensitivity and therefore chosen for the analysis.

to 0.02 rad. For illustration purposes Fig. 9 shows an undersampled grid of this type.

The grid is chosen based on the maximum expected degradation in the detection statistic due to the mismatch between the actual position of a putative source and the template grid, of a few percent. Extensive Monte Carlo simulations were performed to verify the performance of the filter bank. Note that the grid adopted for the search "covers" H1 data parameter space with more redundancy than for L1. The main reason for this difference is the fact that the spanned observation time of the data set used for the L1 detector is longer than that for H1; in fact, the resolution in sky position is highly dependent on the time spanned by the observations. More details are given in Appendix B.

2. Sco X-1

The analysis for Sco X-1 requires a search over the orbital parameters a_{p} and $\bar{T}$ and the gravitational wave frequency f_0. In order to optimally cover this space we consider the metric approach introduced in [48] in the context of binary inspirals and applied to pulsar searches in [10,40]. By using the metric we take advantage of the correlations between the frequency and the orbital parameters—so that a mismatch in orbital parameters can be compensated by a mismatch in frequency—and we therefore reduce the number of orbital templates required to cover the parameter space. We set the total mismatch between signal and template to be 10%. By optimally splitting the mismatch amongst the dimensions of the parameter space, the frequency resolution is set to $\Delta f = 1/(5T_{\mathrm{span}})$ corresponding to $9.285\,051 \times 10^{-6}$ Hz for H1 and $1.066\,382 \times 10^{-5}$ Hz for L1.

The spacing of the filters in the two-dimensional space of the orbital templates is determined primarily by T_{span}—the number of orbital templates in the regime $T_{\mathrm{span}} < P$ scales as T_{span}^5—and the grid orientation is determined by the location of the source in its orbit during the observation. These effects are clearly visible in Fig. 10; although the observation spans differ by only ~15% between L1 and H1, the density of filters for H1 is twice that for L1. The source location within the orbit differs by ≈ 2.22 radians between the L1 and H1 observation periods and correlations between the two orbital parameters are therefore different between the two detectors resulting in template banks that are clearly nonaligned. One further step to optimize the search is to generate separate orbital template banks for each 1 Hz frequency sub-band, because the grid density increases as $(f_0^{\mathrm{max}})^2$, where f_0^{max} is the maximum search frequency. This approach allows an overall gain $\approx$ 30% in computational speed in comparison to using a single template bank with a maximum frequency parameter $f_0^{\mathrm{max}} = 624$ Hz for the whole analysis.

The number of orbital templates used for each 1 Hz subband ranges from 3391 to 3688 in the 464–484 Hz band and from 5738 to 6107 in the 604–624 Hz band for the L1 analysis. The number of frequency filters per 1 Hz band is 93 775; therefore, the number of trials used to cover the parameter space is in the range 3.2×10^8–5.7×10^8. The corresponding numbers for the H1 analysis are 6681–7236, 11 309–12 032, and 108 055, respectively, corresponding

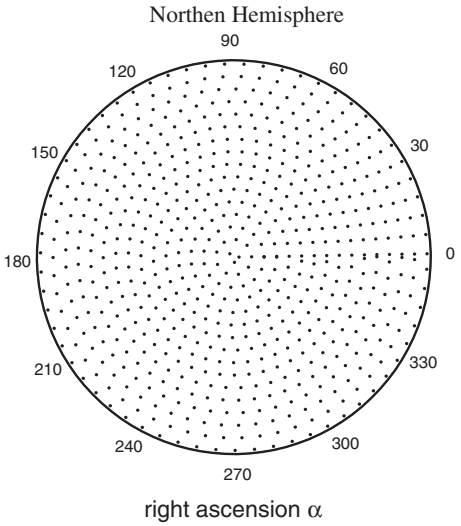

FIG. 9. The sky grid that we have used is of the kind shown here, 25 times more dense. The angular distance between points at the equator in this plot is 0.1 radians. The figure is a projection of the grid on the Northern hemisphere. Distance along the radial direction is proportional to the cosine of the declination.

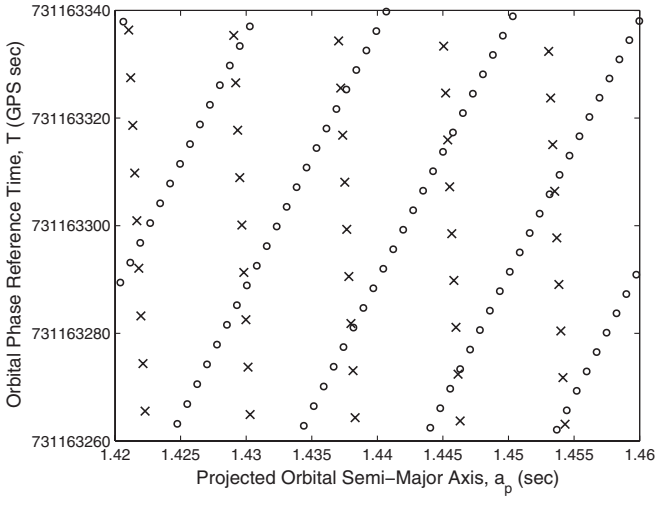

FIG. 10. Here we show a small section ($\approx 1/50$th) of the total orbital parameter space. The crosses and circles represent template locations used to search the L1 and H1 data sets, respectively. This particular template bank was generated for a maximum search frequency of 465 Hz. Note that the templates are *not* uniformly spaced in the $(a_{\mathrm{p}}, \bar{T})$ parameter space, although they appear nearly so in the limited region shown here.

to a total number of trials for each 1 Hz sub-band in the range 7.2×10^8–1.3×10^9.

E. The single detector search

As described in Sec. V B and Fig. 4 the data from each detector is searched over the entire parameter space, by computing the $\mathcal{F}$-statistic for each template in frequency and either position in the sky (for the all-sky search) or orbital parameters (for the Sco X-1 search). In both cases results are stored only for those templates that yield a value of the detection statistic that exceeds the threshold $2\mathcal{F}_{\rm thr} = 20$. This choice is based on limitations on the size of the output files of the search.

Our template banks are highly correlated; thus in order to decrease the number of frequency templates that we store, we treat as correlated the template frequencies which are sufficiently ''close to each other'' and we do not register them as separate templates. The frequency interval that defines how close frequencies have to be in order to be ascribed to the same template is estimated based on the full width at half maximum of the $\mathcal{F}(f_0)$ curve for a representative sample of the parameter space and for random mismatches between signal and template, as would occur in an actual search in the presence of a signal. The resulting frequency intervals are a few times 10^{-4} Hz.

The following information is stored for each template above threshold: the frequency f_0 at which the value of $\mathcal{F}$ is maximum, the values of α and δ for the template, the total width in search frequency bins of the points associated with the maximum, the mean value and the standard deviation of $\mathcal{F}$ over all those points and the value of $2\mathcal{F}$ at the maximum. The same information is stored in the case of the Sco X-1 search, with the orbital parameters a_p and $\bar{T}$ instead of the sky position parameters α and δ.

The computational load for the searches is divided among independent machines, each searching a small frequency region over the entire parameter space. For the isolated pulsar search each CPU analyzes a 60 mHz search band. The processing time for both data streams and the entire sky is typically about 6 hours on a 2 GHz class computer.[5] The typical size of the output, after compression, from a single detector search is around 3 MBytes. For the Sco X-1 search, typically an individual machine searches 0.1 Hz. The equivalent run time on a 60 mHz search band on the entire Sco X-1 orbital parameter space is approximately 9 hours. Although the Sco X-1 orbital templates are fewer than the sky position templates, the two searches are comparable in computational time because the Sco X-1 search uses a greater number of (shorter) SFTs. For the particular data sets selected for this analysis it should be noted that with H1's shorter spanned observation time and fewer SFTs, the computational load is primarily due to the L1 search. The output from the search in total, including both detectors and all search bands, comprises ≈ 1 GByte of results, corresponding to around 700 kBytes per 60 mHz band for the entire orbital parameter space.

Figures 11 and 12 show the distribution of $2\mathcal{F}$ values of the registered templates for sub-bands in reasonably clean spectral regions in both instruments (around 247.1 Hz and 619.5 Hz, respectively) in the top two plots, and in less clean regions in the H1 data (around 329.6 Hz and 465.5 Hz, respectively) in the bottom two plots. In the

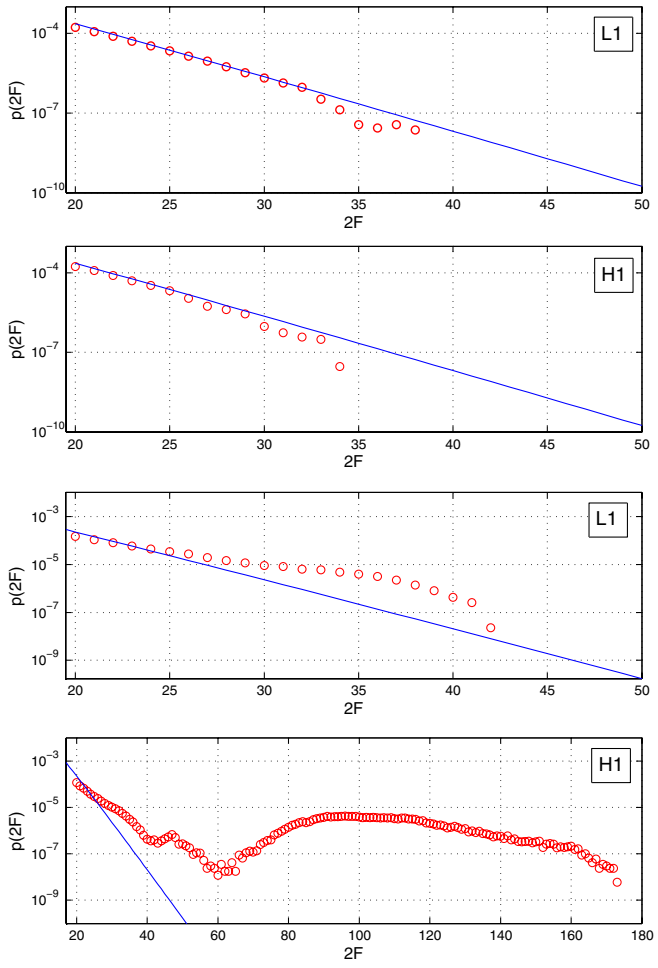

FIG. 11 (color online). The circles show the distribution of $2\mathcal{F}$ values for the templates registered after the single detector all-sky search. The solid line shows the expected distribution for Gaussian stationary white noise. The top two plots refer to the band 247.06–247.12 Hz. The bottom two plots show the same distributions for the 329.56–329.62 Hz band. The expected distribution is dominated for high values of $2\mathcal{F}$ by an exponential term, as is evident from the linear behavior on a semilog scale. In the clean 247 Hz band, the theoretical and the experimental distributions agree very well.

[5]Since the time when this search was carried out, the software has been significantly improved. At the beginning of 2007 the average computing time *per template* of an Einstein@home host performing a coherent search of this type is 0.37 μs $\times$ N_{SFT} , where N_{SFT} is the number of SFTs employed to construct the detection statistic

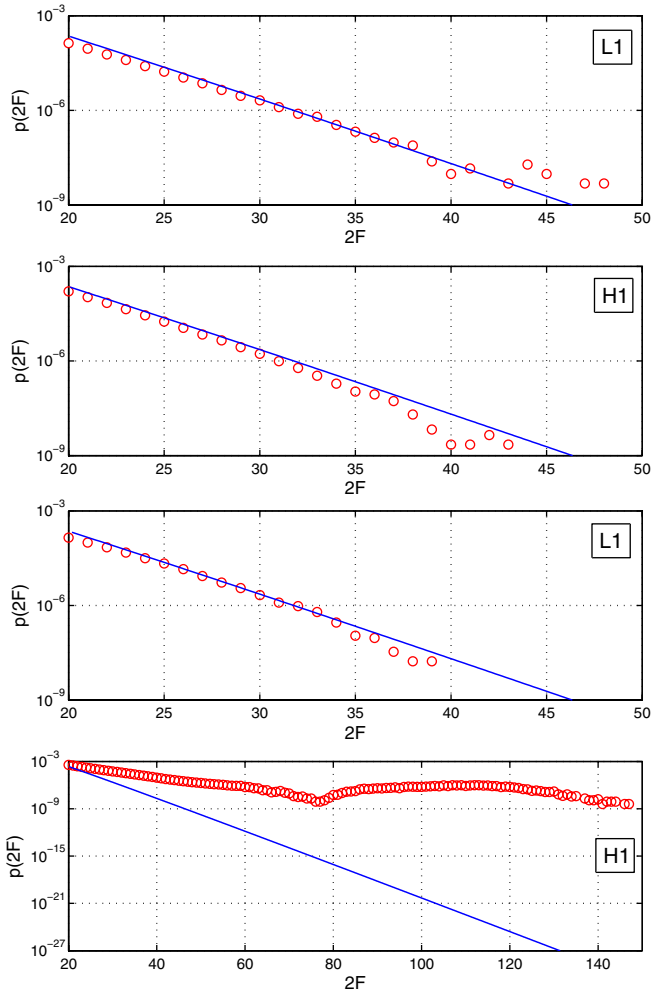

FIG. 12 (color online). The circles show the distributions of $2\mathcal{F}$ values derived from the single detector Sco X-1 searches. The top two plots show the distributions for a clean sub-band in both detectors, 619.0–620.0 Hz. The solid curve represents the theoretical expected distribution. The two bottom plots show the same distributions for the sub-band 465–466 Hz. In this band the H1 results are dominated by large values of $2\mathcal{F}$.

top two plots, the distributions of $2\mathcal{F}$ values closely follow the expected $p_0(2\mathcal{F})$ distribution, Eq. (33). This is not surprising in regions free of evident disturbances, as already shown in [19]. Note that the highest $2\mathcal{F}$ values in the clean bands (top two plots) are higher in Fig. 12 than in Fig. 11. This is due to the fact that the 0.1 Hz Sco X-1 search has more templates than the 0.06 Hz all-sky isolated search.

F. Coincidence analysis

The next stage of the analysis compares the two lists of values of $2\mathcal{F}$ that lie above the threshold $2\mathcal{F}_{\rm thr} = 20$ compiled for each detector. We require that given a template in, say, L1, there exists a template in H1 such that their locations in parameter space are consistent with a physical signal having triggered them both. If this is the case, the relevant values of the detection statistic are stored (the two filters are regarded as "in coincidence"), otherwise they are "rejected" and removed from the lists. This procedure is identical for both searches, but the consistency criteria are different due to the different signals that are searched for. This strategy is effective at reducing the false alarm rate if the noise in the two data streams is uncorrelated. In practice, the data are also populated by a forest of lines present both in L1 and H1, such as 16 Hz harmonics from the data acquisition system and 60 Hz powerline harmonics, and this procedure does not eliminate them. However, it does eliminate the non-Gaussian uncorrelated outliers which also are in the data. We find that the typical sensitivity improvement in h_0 resulting from the coincidence stage is comparable for both searches and in the range 10%–20%, depending on the frequency sub-band.

An additional criterion to identify coincident templates could be based on comparing the values of $2\mathcal{F}$ produced by the two filters; however, as $2\mathcal{F}$ is already maximized over the nuisance parameters ψ and ι, and the integration time of the analyses is shorter than 1 day, it is in practice difficult to introduce an "amplitude consistency cut" that is simultaneously stringent and safe. For this reason we have not included this requirement in the coincidence stage of this search (see, however, the discussion in Sec. VI A).

The coincident templates are then sorted in order of descending *joint significance*. If we indicate with $2\mathcal{F}_{L1}$ and $2\mathcal{F}_{H1}$ the values of the detection statistic for a pair of templates in coincidence, we define their joint significance as

$$s(2\mathcal{F}_{L1}, 2\mathcal{F}_{H1}) = 1 - P_0(2\mathcal{F}_{L1})P_0(2\mathcal{F}_{H1}), \qquad (38)$$

where $P_0(2\mathcal{F})$, defined in Eq. (33), is the single detector false alarm probability for $2\mathcal{F}$, under the assumption that the noise is Gaussian and stationary. We consider the *loudest* coincident template pair as that yielding the largest value of joint significance. In practice, in the numerical implementation we rank events according to $-\{\log[P_0(\mathcal{F}_{\rm L1})] + \log[P_0(\mathcal{F}_{\rm H1})]\}$ with $\log[P_0(\mathcal{F})] = \log(1 + \mathcal{F}) - \mathcal{F}$.

In the remainder of the section we provide details on the specific implementation of the coincidence stage for the two analyses.

1. Isolated neutron stars

For the isolated search the coincidence windows are 1 mHz in frequency f_0 and 0.028 rad angular distance in position on the celestial sphere. These coincidence window values were derived from the results of the Monte Carlo simulations described in Sec. V E. More specifically 0.028 rad represents a mismatch between sky positions of at most 1 grid point. The value of 1 mHz is derived from the

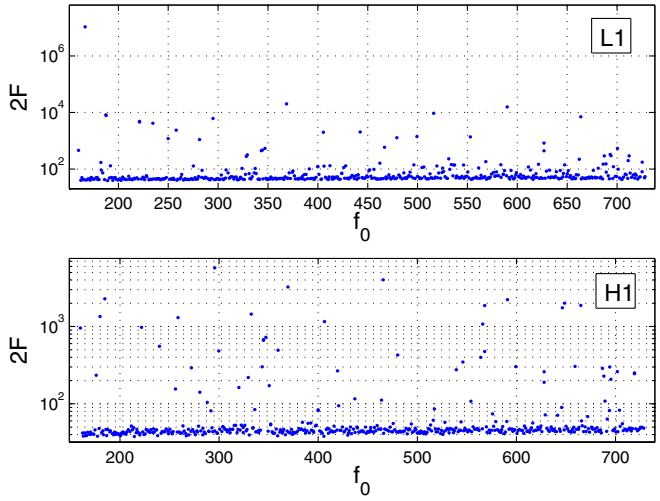

FIG. 13 (color online). These plots show loudest $2\mathcal{F}$ values in the single detector searches in each 1.2 Hz sub-band.

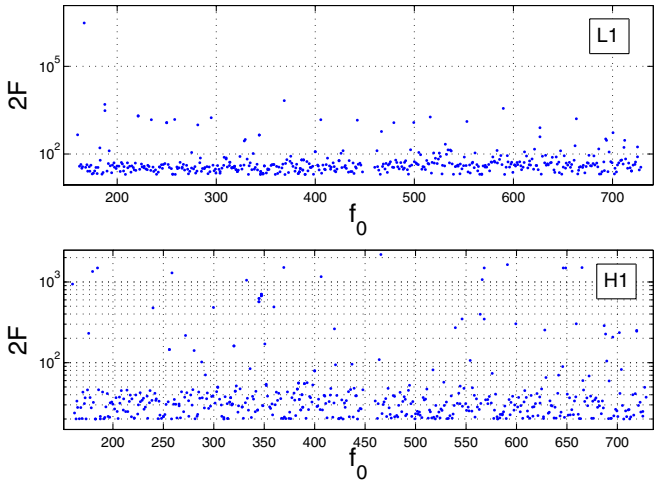

FIG. 14 (color online). These plots show two different views of the values of the detection statistic $2\mathcal{F}$ of the loudest coincident template-couples (one for every 1.2 Hz sub-band) from the isolated pulsar search. The mean value between the threshold and $2\mathcal{F} = 100$ is 39.5 for the L1 data and 32.2 for the H1 data.

results of Monte Carlo simulations by requiring a null false dismissal rate on the simulated sample.

As described in Sec. V D, the all-sky isolated search, if performed on Gaussian white stationary noise, would yield single-interferometer loudest templates in 1.2 Hz sub-bands with mean $2\mathcal{F}$ values of 45.7 for L1 and 41.7 for H1. The difference in these mean values is due to the different data sets used for the two searches (the time spanned by the L1 data set is longer than that of the H1 data set) and by the different location of the detectors on Earth and to the nonuniform antenna pattern of the detectors. In this search, after having excluded outliers with $2\mathcal{F} > 100$, we measure mean values of the loudest templates of 52.2 for L1 and 46.6 for H1.[6] This corresponds to an increased level of spectral contamination in the real data with respect to Gaussian stationary noise. This is not surprising at all—even a simple visual inspection of the spectra reveals that they are contaminated by several ''lines'' (also see the discussion in Sec. VI D in [13]).

After the coincidence step the mean value of $2\mathcal{F}$ for the loudest event is 39.5 for the L1 data and 32.2 for H1. If one compares these values with the mean values *before* coincidence, 52.2 for L1 and 46.6 for H1, one recognizes that the coincidence step yields an improvement in h_0 sensitivity of 15% and 20% for L1 and H1, respectively, [remember from Eq. (35) that $2\mathcal{F} \propto h_0^2$]. In Gaussian stationary noise the expected improvement is 11% and 17% for L1 and H1, respectively. Thus, and again not surprisingly, the coincidence step plays a greater role on real data, which is affected by uncorrelated non-Gaussian disturbances.

Figures 13 and 14 show the values of the detection statistic for the loudest events before and after coincidence, respectively. Figure 15 shows the distributions of $2\mathcal{F}$ for the loudest coincident templates. Figure 16 shows the distribution of loudest coincident templates over the entire sky for all the 1.2 Hz sub-bands. A higher concentration of templates is apparent at the poles. This is to be expected since the poles are the regions from where a monochromatic signal would be received by our detectors at a nearly constant frequency. In other words, spectral artifacts at

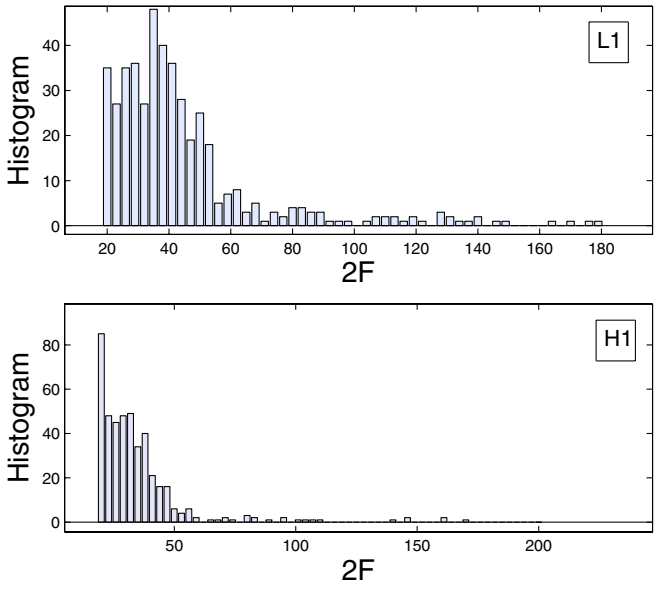

FIG. 15 (color online). These plots show the distributions of the values plotted in Fig. 14. The mean value of these distributions depends on the volume of the parameter space that the search extends over. In this case it is the whole sky in 1.2 Hz frequency sub-bands.

[6]The $2\mathcal{F} > 100$ cut has been made only when computing the mean values reported above in order to eliminate large outliers that would have dominated the mean; see Figs. 14 and 15.

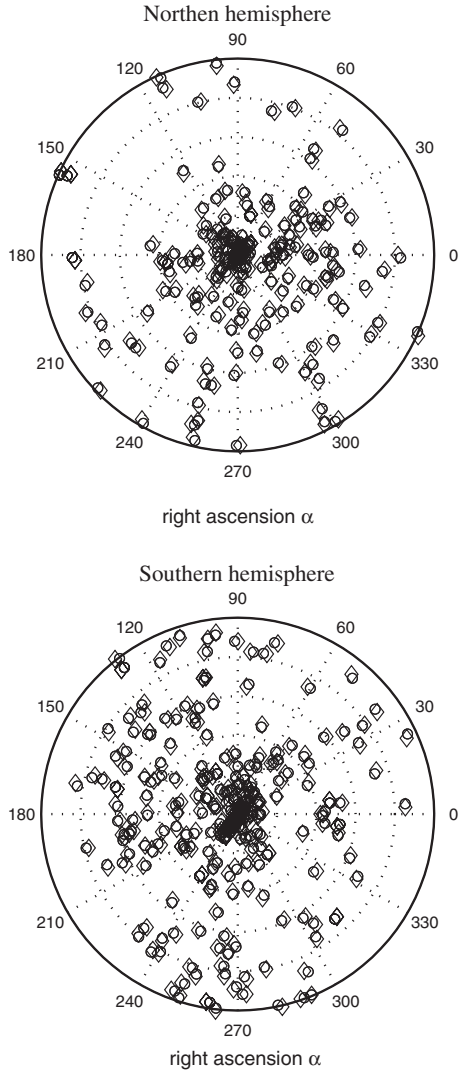

FIG. 16. These figures show the location in the sky of the loudest coincident template in each 1.2 Hz sub-band. The circles show the templates found with the L1 search and the diamonds show the coincident templates found with the H1 search.

fixed frequency are consistent with sources close to the poles, during our observation time.

2. Sco X-1

In the Sco X-1 analysis we identify coincident templates in L1 and H1 by using the metric on the relevant parameter space; see Sec. V D 2. A real signal, if present and of sufficient amplitude, will trigger templates in both detectors. These templates will be in close proximity in parameter space but not necessarily identical. The procedure starts by taking each template above threshold from the L1 detector and first testing for orbital parameter consistency with templates from the H1 detector. We do this using a property inherent to the orbital template bank, which is schematically illustrated in Fig. 17. We know that a signal will return at least 90% of its optimal detection statistic

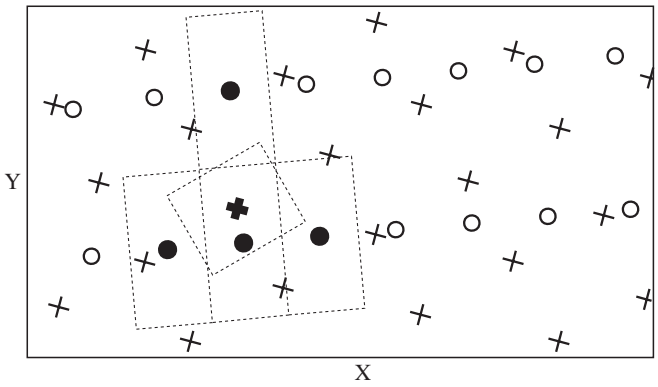

FIG. 17. Here we represent a small region of orbital parameter space. The crosses represent orbital templates used in the L1 search and the circles represent those used in the H1 search. Note that the template orientation and spacing are not the same in each template bank. The bold × represents the location of a candidate event in the L1 detector, and the dashed rectangle surrounding it represents the area within which a signal must lie if this template is the closest to the signal. The filled circles represent those templates in the H1 detector that are possible coincident candidates. They are identified due to the overlap between their respective dashed rectangular regions and the L1 candidate event dashed region. In the S2 search there are on average 12 orbital templates in H1 that are consistent with each template in L1.

when processed using the "closest" template (in the absence of noise), where the distance between templates is defined by the metric. We can therefore identify a rectangular region around the L1 filter in the two-dimensional plane of orbital parameters. This region is easily calculated from the metric used to place the templates, and we would expect the true signal parameters to lie within it. We can now repeat this process for each H1 template and construct, based on the metric associated with H1, the region covered by the H1 filters (i.e. the filters that are associated with a value of $2\mathcal{F}$ above threshold). The test now becomes a simple matter of checking for any overlap between the region covered by the L1 filter under scrutiny and the H1 filters. Overlap implies a possible consistent signal location (in orbital parameters) capable of triggering a filter in L1 and H1. In this analysis there are on average 12 orbital templates in H1 that are "consistent" with each template in L1. The process that we have just described is then repeated for all the L1 filters.

So far we have considered only the orbital parameters. The second stage to identify filters in "coincidence" is to test for frequency consistency amongst those filters that have survived the previous test. The use of the projected metric described in Sec. V D 2 exploits the correlations between the orbital templates and the gravitational wave frequency to reduce the overall number of filters. Doing so allows greater differences between the true and detected source orbital parameters and greater differences between the true and detected gravitational wave frequency. Using Monte Carlo simulations we measured the maximum sepa-

ration between a signal's true and detected frequency. This separation is largely determined by the spanned observation time. For the data sets chosen, the maximum separation for L1 was found to be 2.158×10^{-4} Hz and for H1 was 1.773×10^{-4} Hz. This corresponds to a maximum separation of 3.931×10^{-4} Hz between a candidate in L1 and H1 in order to be consistent with a common signal. This is the frequency coincidence window that we have chosen. Note that it is equivalent to ± 40 frequency bins in the H1 search.

If a pair of candidate events is found to be consistent in both orbital parameter space and frequency space then they are classed as a coincident event. Note that a single candidate event in the L1 detector can have many possible coincident pairs in the H1 detector (and vice versa).

The power of the coincidence analysis is shown in Fig. 18 where the effect of the coincidence constraint is seen to reduce the values of our loudest events. Before coincidence the average value of $2\mathcal{F}$ for the loudest events (excluding three 1 Hz sub-bands that contain major spectral disturbances: 465–466 Hz and 479–481 Hz) was 40.8 for L1 and 45.4 for H1. After coincidence this becomes 28.6 for L1 and 33.5 for H1 which corresponds to an improvement in h_0 sensitivity of $\sim$16%. This is broadly consistent with the results obtained for the isolated neutron star analysis. In Fig. 19 we show the location of the coincident templates in the orbital parameter plane that produce the loudest event in each of the 40 1-Hz sub-bands.

G. Upper limits

We place an upper limit on the amplitude of the gravitational wave signal (from either a population of isolated neutron stars or from Sco X-1) in every $\sim$1 Hz sub-band on which the search was performed; the upper limit is based on the loudest coincident event found in that band during the search. The procedure employed is conceptually identical to the one used in [19] to set an upper limit on the emission from J1939 + 2134, given the measured values of

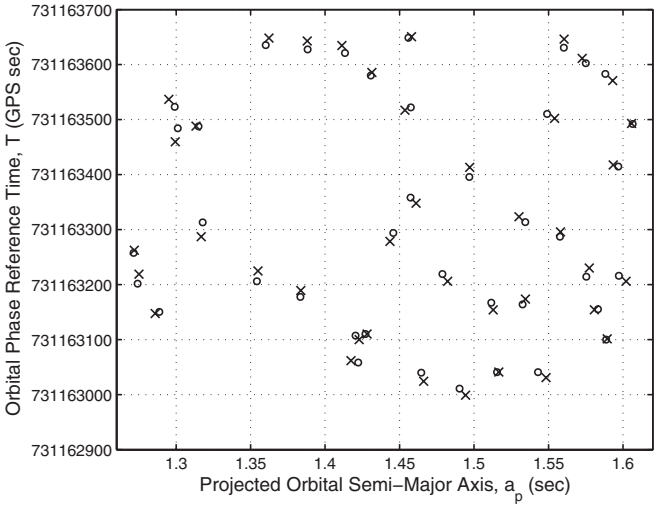

FIG. 19. The locations in orbital parameter space of the loudest events found in each 1 Hz sub-band. The crosses represent events found in L1 and the circles represent events found in H1.

the $\mathcal{F}$ statistic for that targeted search. In this section we describe the Monte Carlo procedure; further details are provided in Appendix D.

Let $s^*(f_0)$ indicate the measured value of the joint significance, see Eq. (38), of the loudest coincident event in the sub-band beginning at frequency f_0. For every sub-band a set of N (typically several thousands, see Appendix D for more details) injections of fake signals in the real data is performed at fixed amplitude h_0. Each injection is searched for in the data and, if detected as a coincident event, its significance is computed. A confidence level $C(h_0)$ is assigned to this set of injections

$$C(h_0) = n(h_0)/N \tag{39}$$

with $n(h_0)$ being the number of trials out of N in which the measured joint significance of the injected signal is greater

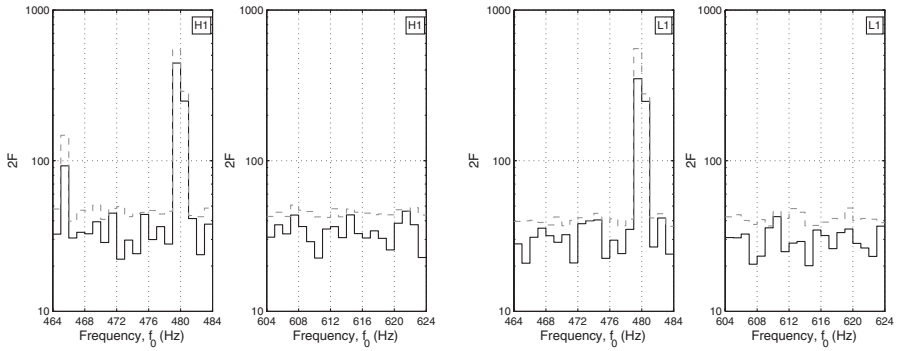

FIG. 18. Here we show the effect that the coincidence analysis has on the loudest measured detection statistic in each 1 Hz sub-band within the Sco X-1 parameter space. The solid black curves represent the loudest coincident $2\mathcal{F}$ values. The dashed gray curves represent the $2\mathcal{F}$ values before the coincidence analysis. Note that in clean sub-bands there is a reduction of $\sim$1.4 in the loudest $2\mathcal{F}$ value.

than or equal to s^*. Equation (39) defines the h_0 upper limit value as a function of the confidence C.

For every injection in a set at fixed h_0 the remaining parameters are chosen randomly from within the boundaries of the relevant search parameter space. The errors in the estimate of the upper limit are directly computed from the Monte Carlo results. Our approach typically yields uncertainties in the values of the upper limit of $\sim$1%–3%. An additional and larger uncertainty arises from the instrument calibration, which varies with time and depends on the detector and the frequency band. Over the entire parameter space, we estimate that for the data sets used in the analyses presented here, the uncertainties amount to 11% and 9% for the isolated neutron star and Sco X-1 analysis, respectively. These estimates are conservative.

The tables in [49] detail all the upper limit results. The uncertainties associated with the upper limit Monte Carlo procedure are reported separately from the calibration uncertainties and typically they are smaller than the latter.

VI. RESULTS

In this section we present the results of the analysis performed using the pipeline shown in Fig. 4 and described in Sec. V. We first discuss the results regarding the all-sky search for isolated neutron stars and then turn to upper limits on radiation from Sco X-1.

A. Isolated neutron stars

Figure 20 shows the 95% upper limits on h_0 for every 1.2 Hz wide sub-band over the whole sky. The values of the frequency refer to the lower extremum of each sub-band. (The upper limit values, along with their estimated uncertainties, may also be found in tabular form in [49].) The circles around the upper limit dots mark points in the 90th percentile in joint significance. About 2/3 of these points are also in the 90th percentile for $h_0^{95\%}$.

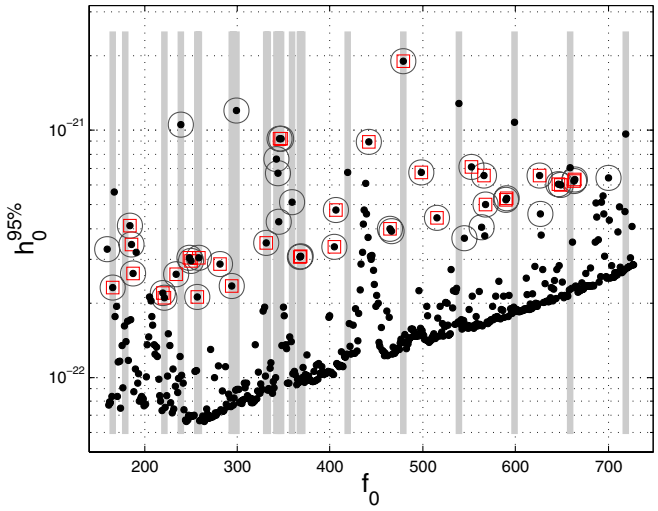

FIG. 20 (color online). Upper limits based on the loudest template over the whole sky in 1.2 Hz sub-bands. The vertical stripes mark the sub-bands containing known spectral disturbances. The circles mark the 90th percentile most significant results. The squares indicate that the values of the detection statistic in the two detectors are not consistent with what one would expect from an astrophysical signal.

About one quarter of the 90th percentile significance points lie in sub-bands influenced by spectral disturbances (points with circles on shaded bands in Fig. 20). Most of the remaining points can be immediately attributed to nonastrophysical sources because the ratio of the $\mathcal{F}$-statistic values in the two detectors is either too large or too small to be consistent with being due to the same signal. These points are indicated in Fig. 20 by a square. There remain 6 points which are in the 90th percentile in significance and cannot be excluded based on the ratios of the $\mathcal{F}$-statistic values. They appear at the frequencies 160.0 Hz, 466.79 Hz, 546.03 Hz, 564.02 Hz, 626.80 Hz, and 700.51 Hz. The 160.0 Hz frequency coincides with the 10th harmonic of 16.0 Hz, a key operating frequency of the data acquisition system. We are thus confident that the origin of this outlier is instrumental. The points at 466.79 Hz, 626.80 Hz, and 700.51 Hz are due to lines only in the L1 instrument which have disappeared in science runs subsequent to S2. This check suggests that the outliers are of instrumental origin. The 546.03 Hz and 564.02 Hz points are due to lines which clearly appear only in H1. However, the lines are present in the S2 run and in later science runs. The amplitude of both lines *decreases* with increasing sensitivity of the instrument, dropping by a factor of 10 (in noise power) as the sensitivity increases by a factor greater than 5. This indicates a behavior which is not consistent with the model of the signal that we are searching for here and suggests that these lines are of instrumental origin. Figure 21 shows the average power spectral density in both detectors in the frequency regions where these outliers are located.

Unlike what is described in [13] no frequency band is excluded from the upper limit analysis due to it being contaminated by known noise artifacts. This results in extremely loud events in some sub-bands: those containing the 60 Hz powerline harmonics, the L1 calibration line (at 167 Hz), the violin modes of the suspension wires of the test mass (in the 340–350 Hz region) and the various oscillator harmonics at multiples of 36 Hz together with the beating of the 0.74 Hz pendulum mode of a test mass against the oscillator line (in the 220–335 Hz region). In the case of the 179.4 Hz sub-band containing the 180 Hz powerline harmonic, the spectral disturbance is so strong that the upper limit Monte Carlo does not converge to an upper limit $h_0^{95\%}$ value. The sub-bands marked by shaded vertical stripes indicate frequencies where known spectral artifacts are present.

The upper limit values presented here are in broad agreement with what is expected and we consider this a

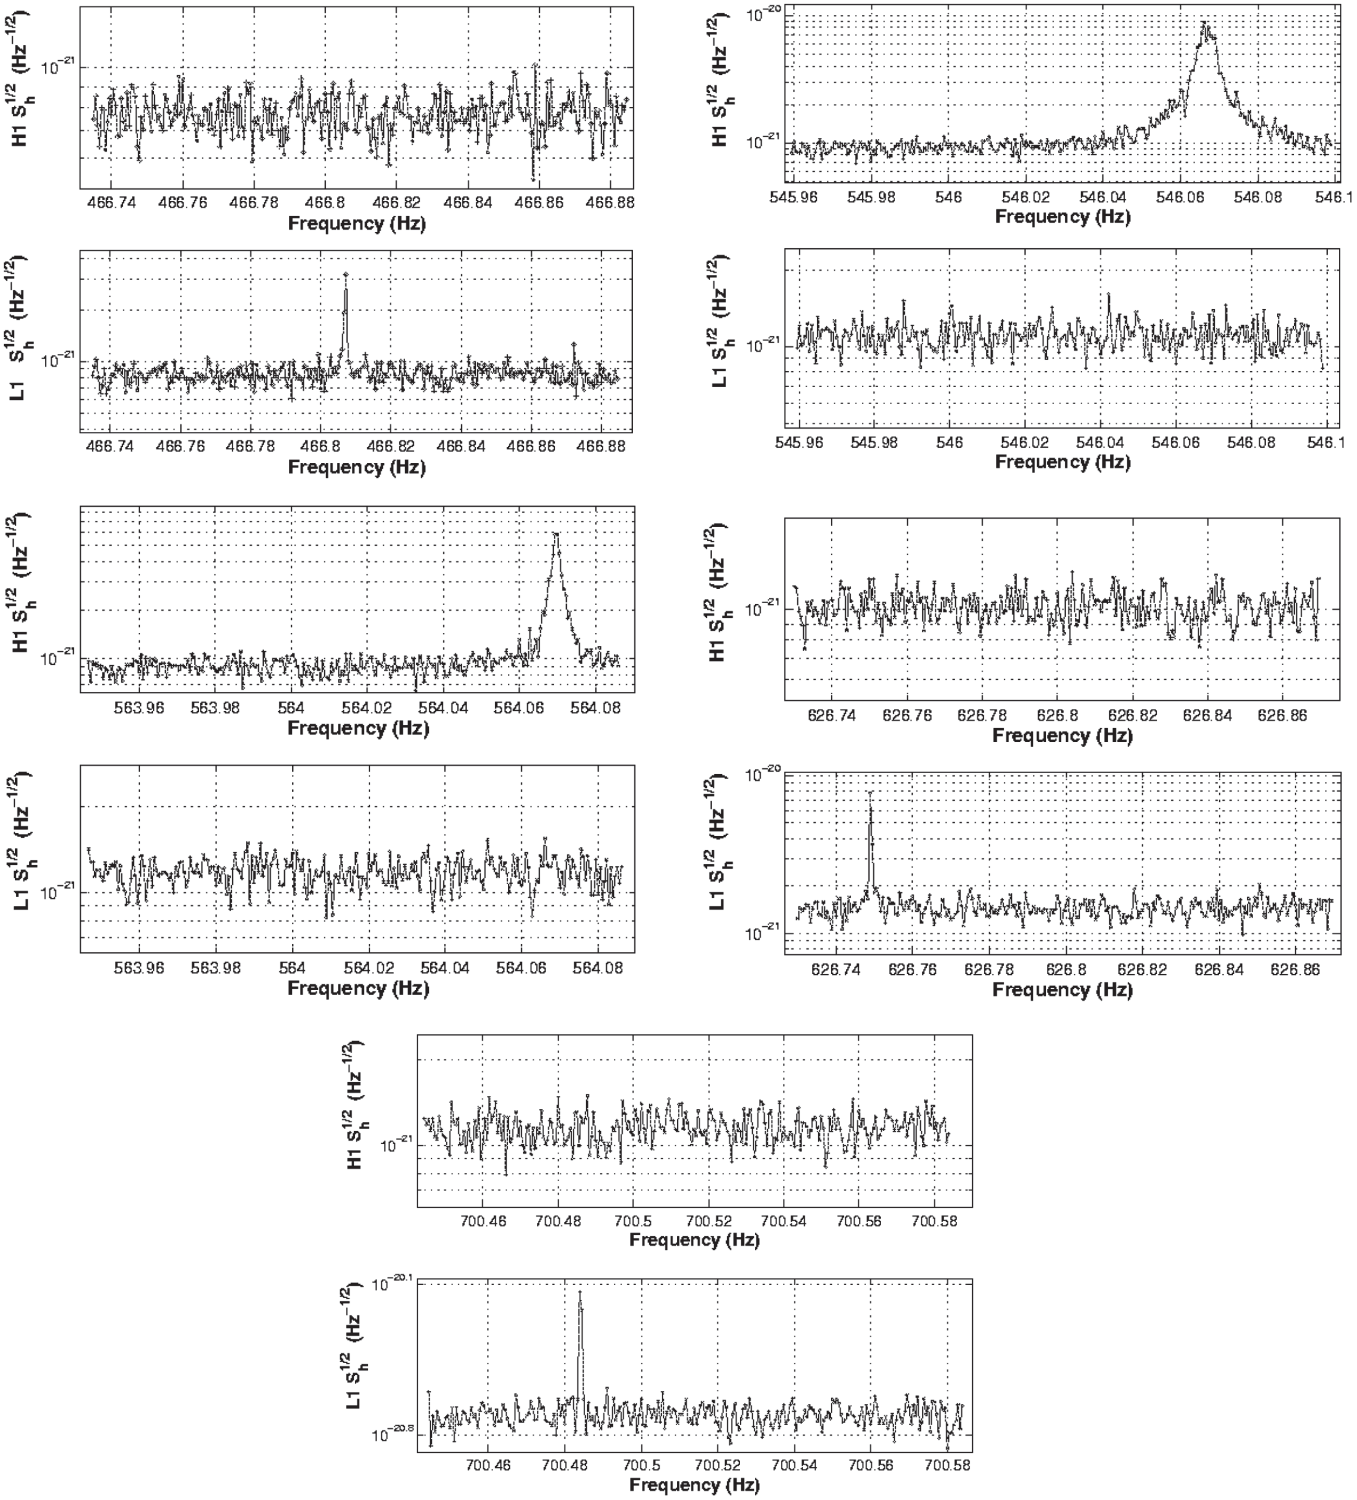

FIG. 21. The average power spectral density in the two detectors in the frequency regions around the five not immediately explained outliers from the search for isolated pulsars. The width of these disturbances is sufficiently small that they could not be discarded as of nonastrophysical origin based on this. A peak in power spectral density is clearly visible only in one of the detectors, but the measured values of the detection statistics are not inconsistent with an astrophysical signal, albeit a rare one. As explained in Sec. VI A we very strongly suspect these excesses of power are not due to a continuous wave source because of their inconsistent amplitudes in science runs subsequent to S2.

further validation of the analysis pipeline. We have run the pipeline presented here on Gaussian stationary noise and empirically derived a formula for the expected $h_0^{95\%}$ as a function of the noise level in the detectors:

$$h_0^{95\%} \sim 29.5\sqrt{\frac{\langle S_h \rangle}{10\ \text{hrs}}}, \tag{40}$$

with $\langle S_h \rangle$ being the average noise level over 1.2 Hz and over the observation time in every detector and then averaged over the detectors: this quantity is shown in Fig. 6. We would like to stress that Eq. (40) refers to this particular analysis and pipeline. The expected upper limits for Gaussian stationary noise are plotted against the measured ones in Fig. 22. It is clear that in regions where the data is not Gaussian and stationary Eq. (40) does not predict correctly the values that we measure and the discrepancy between the prediction and the measured value depends on the details of the spectral disturbance and of the method used for estimating the noise. This is particularly evident close to spectral disturbances, where clearly the noise is not white Gaussian and often not stationary and the predictions can even be larger than the actually measured upper limit value (see the points below 1 in the lower plot of Fig. 22). However, the ratio of the measured upper limits to the expected one never exceeds 4.4 and the 90th percentile level in this ratio is 1.7.

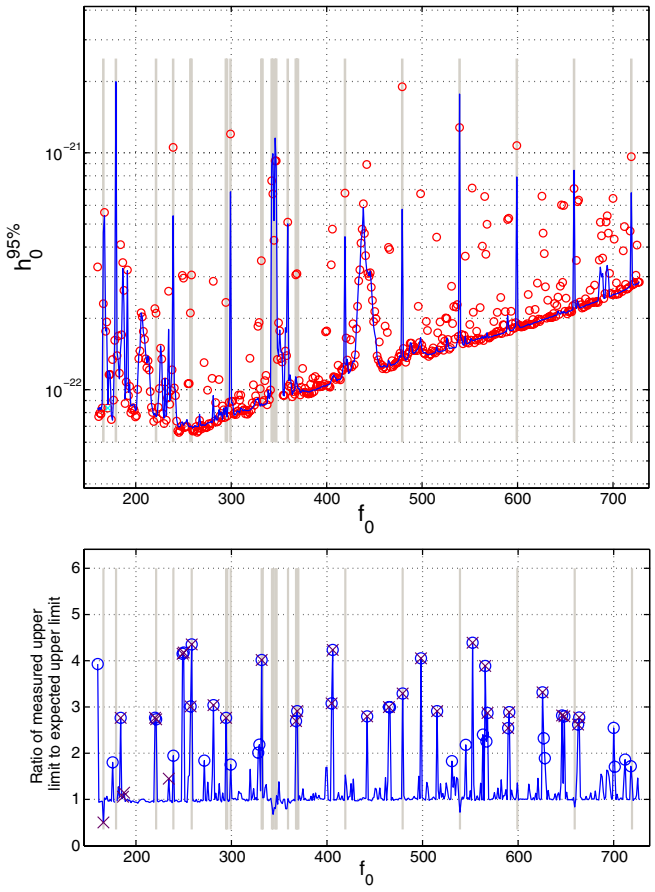

FIG. 22 (color online). Top plot: the circles are the upper limits based on the loudest template over the whole sky in 1.2 Hz sub-bands. The solid line is the expected upper limit under the assumption that the noise in the detectors is Gaussian and stationary. Bottom plot: ratio of the measured upper limit values to the expected upper limit values. The circles indicate the 90th percentile values of the measured to expected upper limit ratio. The crosses mark the upper limit values for which the ratio of the detection statistic values in the two detectors is not consistent with what we would expect from a signal. In both plots the shaded regions indicate frequency bands affected by known spectral disturbances.

B. Sco X-1

The upper limits on gravitational waves from Sco X-1 are summarized in Fig. 23 (more details are provided in [49]): we show $h_0^{95\%}$ over 1 Hz sub-bands in the range 464–484 Hz and 604–624 Hz, assuming that the source is in an exactly circular orbit. We would like to stress that these limits apply to a source whose orbital parameters lie in the region reported in Table I, corresponding to $1-\sigma$ errors. The typical value of $h_0^{95\%}$ is $\approx 2\times10^{-22}$ over the whole analyzed 40 Hz band, with the exception of a band ≈ 2 Hz around 480 Hz which corresponds to one of the strong harmonics of the 60 Hz powerline, cf. Fig. 7. In this region the upper limit is $h_0^{95\%} \approx 10^{-21}$. Such values are consistent with the sensitivity estimates shown in Fig. 3, which were derived under the assumption of Gaussian and stationary noise and include a number of approximations to quantify the effects of each stage of the pipeline considered in this

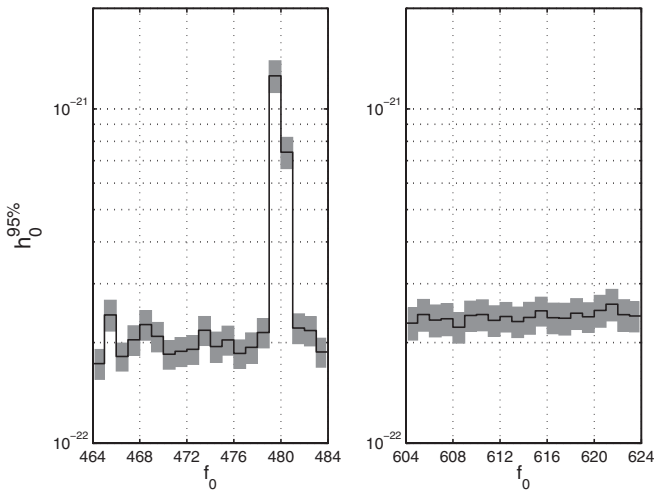

FIG. 23. The upper limits on the amplitude of gravitational wave radiation from Sco X-1. The plot shows the limits on h_0 at 95% confidence (solid line) as a function of frequency. We report one limit for every 1 Hz sub-band and the Sco X-1 orbit is assumed to be exactly circular. The shaded region shows the combined errors on $h_0^{95\%}$ due to the injection process (typically in the range $\pm 2-5\times10^{-24}$) and instrument calibration. The latter, estimated at the level $\approx 9\%$, dominates the uncertainties.

search. Through the statistical modelling of the pipeline we are able to express the expected $h_0^{95\%}$ upper limit as

$$h_0^{95\%}(f) \sim 28\sqrt{\frac{\langle S_h(f)\rangle}{6\ \text{hrs}}}, \tag{41}$$

where $\langle S_h(f)\rangle$ is the noise level over the 1 Hz sub-bands averaged over the observation time, the frequency band, and the detectors.

In Fig. 24 we show the value of h_0^C as a function of C for selected frequency sub-bands. We have considered both "quiet" and "noisy" spectral intervals, but have restricted this analysis to only four 1 Hz frequency sub-bands due to computational burdens. Figure 24 shows that $h_0^{99\%}$ and $h_0^{50\%}$ are a factor ≈ 2 larger and smaller, respectively, than $h_0^{95\%}$.

So far we have restricted the discussion of the upper limits to the case of an exactly circular orbit. This is the model that we have assumed in building the templates used in the analysis. As we have mentioned in Sec. IV B 2, the orbital fits of the optical data are consistent with $e = 0$, which is in agreement with the theoretical expectations inferred from evolutionary models. However, present observations do not constrain the eccentricity to $e \lesssim 10^{-4}$, which would introduce *systematic* losses of signal-to-noise ratio smaller than ≈ 0.1, the value of the mismatch adopted for this search. It is therefore important to explore the consequences of a (unlikely but possible) nonzero eccentricity of the Sco X-1 orbit on the results reported so far. The pipeline that we have developed allows us to quantify this effect in a fairly straightforward way: the Monte Carlo software injections used to set upper limits are performed again by drawing signals from a population of binaries where now the eccentricity is set to a (constant) value $e \neq 0$. The orbital parameters and the frequency are chosen randomly exactly as in the case for a circular orbit. We detect the signals from eccentric orbits with the search pipeline constructed with a bank of filters for a perfectly circular orbital model. In this way we can quote consistently an upper limit on h_0 for $e \neq 0$. We repeated this procedure for selected values of the eccentricity, $e = 10^{-4}$, 5×10^{-4}, 10^{-3}, and 5×10^{-3}. The dependence of the value of the upper limit on h_0 as a function of the confidence level for four representative frequency sub-bands is shown in Fig. 24 and the upper limits over the whole 40 Hz region (at fixed confidence level) are summarized in Fig. 25 and in [49]. Notice that in this case we choose different values of the confidence level depending on the eccentricity of the orbit of the putative source population used for the injection. This stems directly from the fact that the detection efficiency of the search pipeline is progressively reduced as the model of the injected signals differs more and more from that of the detection templates. In other words we suffer from systematic losses of signal-to-noise ratio due to the fact that the templates are not properly matched to the signal: for $e \gtrsim 10^{-3}$ the fitting

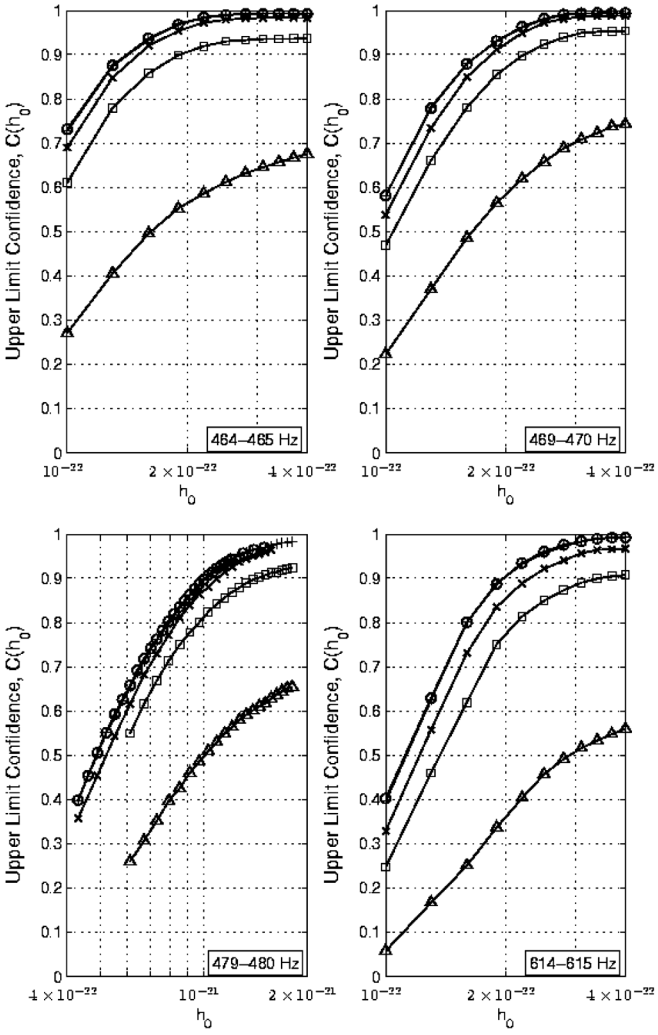

FIG. 24. Examples of how the confidence level in the upper limit on h_0 from Sco X-1 scales with the amplitude of injected signals. Here we show 4 plots each corresponding to a different 1 Hz sub-band. The sub-bands 464–465 Hz, 469–470 Hz, and 614–615 Hz show the confidence level across the range of typical injected h_0 values (10^{-22}–4×10^{-22}). The sub-band 479–480 Hz contains a large spectral disturbance and the injections are appropriately increased in amplitude in order to achieve the required confidence level. The 5 curves in each plot represent the confidence level associated with different values of orbital eccentricity. The circles represent $e = 0$, the plusses represent $e = 10^{-4}$, the crosses represent $e = 5 \times 10^{-4}$, the squares represent $e = 10^{-3}$, and the triangles represent $e = 5 \times 10^{-3}$. Note that for the 6-hour observation, the pipeline is as sensitive to signals with values of $e < 10^{-4}$ as it is to circular orbits. For signals with $e = 10^{-3}$ the pipeline has a maximum detection efficiency of ~90%–95%. For $e = 5 \times 10^{-3}$ the maximum detection efficiency reaches only ~50%–70%.

factor of a filter generated by modelling Sco X-1 as a circular orbit binary is <0.9. Indeed, regardless of the strength of the injected signals, the pipeline is unable to detect at least 95% of them, see Fig. 24. We find that for $e = 10^{-3}$ and 5×10^{-3} the pipeline has a maximum de-

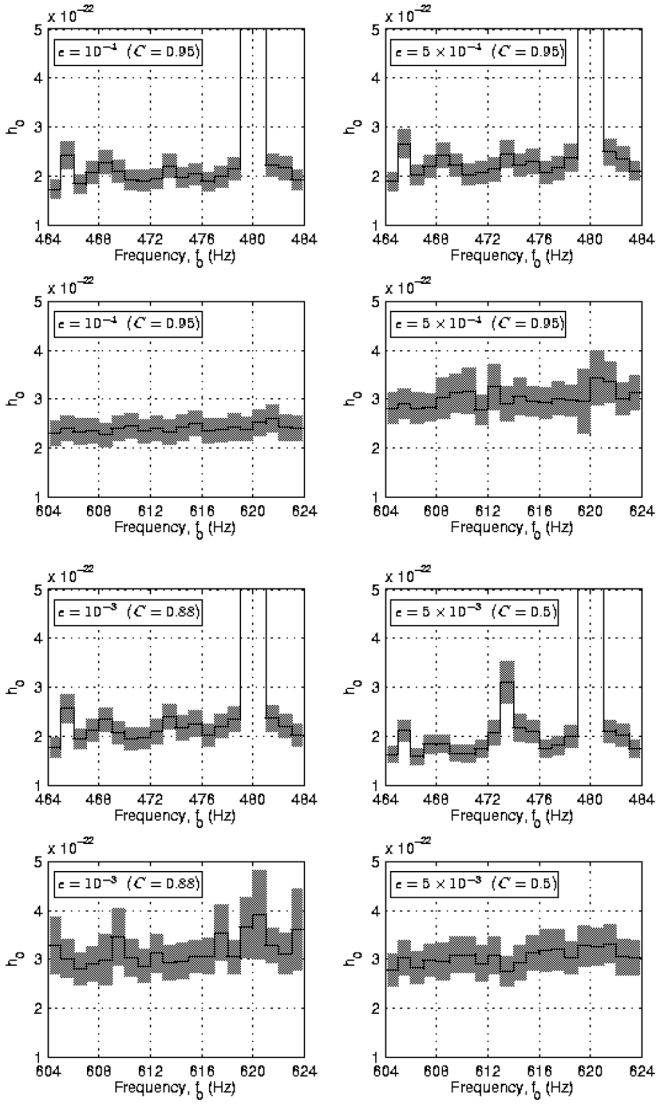

FIG. 25. The upper limits on h_0 as a function of frequency for different orbital eccentricities for the Sco X-1 search. Notice that for sake of clarity, we have adopted a different scale on the vertical axis than that used in Fig. 23; as a consequence the upper limits in the frequency region 479–481 Hz are not shown because they are off scale. Because of the systematic loss of signal-to-noise ratio due to signal-template mismatch for populations of signals from sources in eccentric orbits, the confidence level C at which the upper limit h_0^C is computed is different depending on the value of eccentricity e; see text for a detailed justification. We report $h_0^{95\%}$ for $e = 10^{-4}$, $h_0^{95\%}$ for $e = 5 \times 10^{-4}$, $h_0^{88\%}$ for $e = 10^{-3}$, and $h_0^{50\%}$ for $e = 5 \times 10^{-3}$. The shaded areas represent the combined errors due to the injection process and the instrument calibration.

tection efficiency in the ranges ~90%–95% and ~50%–70%, respectively. As a consequence, for $e = 10^{-3}$ we report $h_0^{88\%}$, and for $e = 5 \times 10^{-3}$ we consider $h_0^{50\%}$, because across each of the 1 Hz sub-bands we have achieved at worst an 88% and 50% confidence level, respectively (see Fig. 25). On the other hand, for $e \leq 5 \times 10^{-4}$ the systematic loss of signal-to-noise ratio is small or even negligible, and we can quote upper limits at 95% confidence level. We find that, as expected, the values of $h_0^{95\%}$ for $e = 0$ and $e = 10^{-4}$ are essentially identical. For $e = 5 \times 10^{-4}$, $h_0^{95\%}$ is about 50% bigger than in the case $e = 0$.

Considering the limited sensitivity of the present analysis (see Figs. 3, 23, and 25) with respect to the astrophysical predictions, Eq. (24), we have not followed up (e.g. using a longer integration time T_{span}) regions of the parameter space that yielded particularly large values of $\mathcal{F}$. Such a follow-up would be computationally very intensive and the fact that we are targeting a continuous gravitational wave emitter allows us to go back to the same parameter space in the future, exploiting higher sensitivity, better quality data, and a more sensitive search algorithm. This work is already in progress. It is, however, important to establish that the results that we have obtained do not show any obvious unexplained feature and are qualitatively consistent with the expectation that no signal is present in the data set at the sensitivity level of the search. In Fig. 26 we show the distribution of the parameters that characterize the filters in coincidence in L1 and H1 for a representative frequency band, 614–614.2 Hz. Because of the high correlation of the templates used in the analysis, one would expect a cluster in parameter space of filters in coincidence, were a real signal present. Considering the coincident filters in the three-dimensional search space $(f_0, a_{\mathrm{p}}, \bar{T})$ and projecting them onto the plane $(a_{\mathrm{p}}, \bar{T})$, no particular structure is evident, with coincident templates evenly distributed across the plane. This is also broadly consistent with the distributions of coincident templates that we have obtained by performing Monte Carlo simulations of the entire search pipeline on stationary and Gaussian noise. There is, however, some structure in the plane $(a_{\mathrm{p}}, \bar{T})$ that is determined by narrow spectral disturbances and accounts for the large outliers in the values of the detection statistic at the end of the analysis pipeline. In order to explore this, it is useful to project the same three-dimensional parameter space onto either the (f_0, a_{p}) or $(f_0, \bar{T})$ plane. Structures are now clearly visible consisting of "stripes" of events at approximately (but not exactly) constant frequency. These are caused by small narrow spectral features present in the data that produce relatively large values of the $\mathcal{F}$-statistic for a number of orbital templates. Because of the short coherent integration time, such disturbances are not averaged out by the demodulation process and are registered in the single detector search. We find that they are very common in the output of the single detector search; the coincidence stage of the analysis allows those stripes of events that exist in both detectors at approximately the same frequency to survive the entire pipeline.

We have so far reported the results of the analysis as upper limits on the signal amplitude h_0. We can now recast

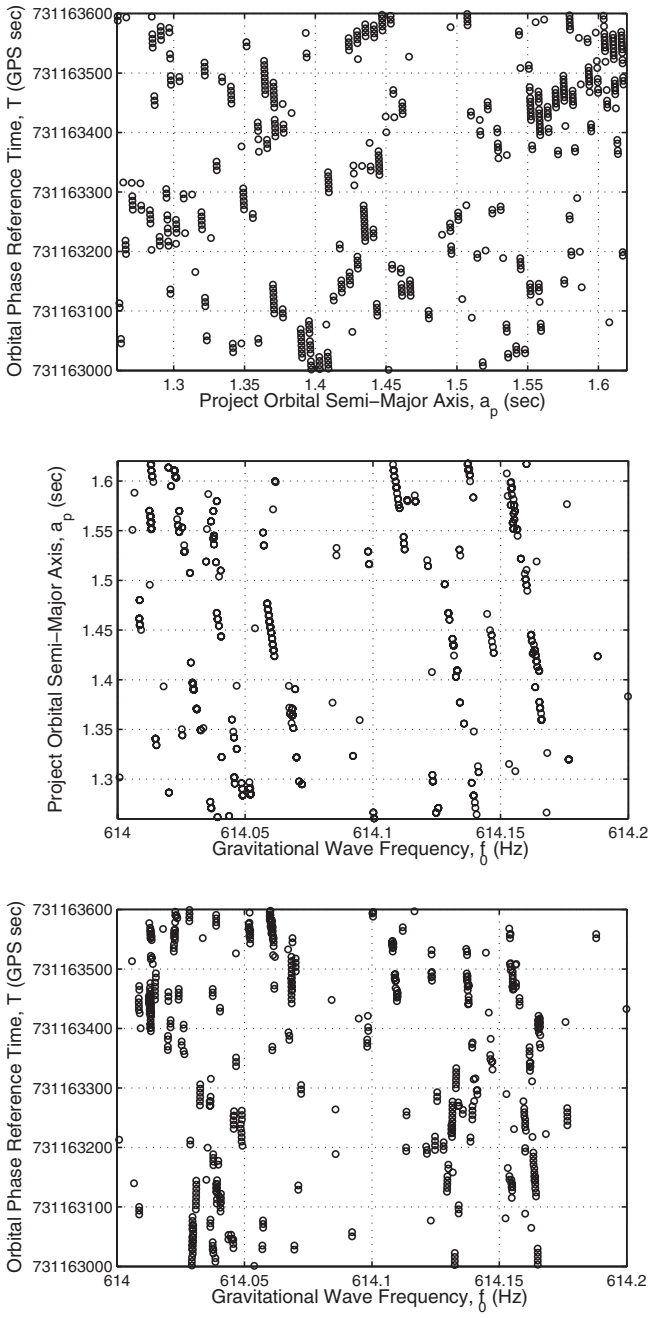

FIG. 26. The locations of coincident templates within a representative subsection of the Sco X-1 parameter space. These plots contain events found in the frequency band 614.0–614.2 Hz and are representative of a ''clean'' search band. We show only events from the H1 detector for clarity (L1 coincident events lie in approximately the same locations). We show 3 two-dimensional projections through the three-dimensional search space. The first plot shows the projected orbital semimajor axis, a_{p}, versus the orbital phase reference time, $\bar{T}$, and shows no obvious structure. The second and third plots show a_{p} and $\bar{T}$, respectively, versus f_0. Here we do see structure caused by small narrow lines present in the data that the demodulation process has failed to smooth out over the relatively short 6-hour observation time.

them as upper limits on the ellipticity ϵ of the neutron star, taking the distance of Sco X-1 as $d = 2.8$ kpc; see Table I. Using

$$\epsilon \simeq 0.237\left(\frac{h_0}{10^{-24}}\right)\left(\frac{d}{1\ \mathrm{kpc}}\right)\left(\frac{1\ \mathrm{Hz}}{f}\right)^2\left(\frac{10^{45}\ \mathrm{g\,cm^2}}{I_{zz}}\right), \quad (42)$$

and the canonical value for the principal moment of inertia $I_{zz} = 10^{45}$ g cm^2 we obtain:

$$\epsilon^{95\%} = 4.0 \times 10^{-4}\text{–}3.6 \times 10^{-3} \quad \text{for } e \leq 10^{-4}, \quad (43)$$

over the frequency band. The previous result can be generalized to the case of a more pronounced nonzero eccentricity; e.g. we obtain

$$\epsilon^{88\%} = 5.1 \times 10^{-4}\text{–}3.7 \times 10^{-3} \quad \text{for } e = 10^{-3}. \quad (44)$$

Despite being far from astrophysically interesting, Eqs. (43) and (44) represent the first direct measurements of the ellipticity of the neutron star in Sco X-1 in the relevant frequency band.

VII. CONCLUSIONS

We have presented here results from two wide parameter space searches for continuous gravitational wave signals. A subset of data from the second science run of the LIGO instruments was analyzed, the data chosen to maximize the sensitivity of the search. Two different astrophysical searches were performed: an all-sky search aimed at signals from isolated neutron stars and an orbital parameter search aimed at signals from the neutron star in the binary system Sco X-1. Both searches also cover a wide range of possible emission frequencies: a 568.8-Hz band for the isolated pulsars search and two 20-Hz bands for the Sco X-1 search.

The sensitivity of these analyses makes the detection of a signal extremely unlikely. As a consequence the main goal of the paper is to demonstrate an analysis method using real data, with a pipeline considerably more complex than any other coherent searches previously performed. More importantly, this coherent search will be deployed in a hierarchical analysis scheme. The first step of a hierarchical analysis sets the ultimate sensitivity of the search: candidates that do not survive the first threshold are lost. It is thus crucial to employ the most sensitive possible technique in this first step. The coherent search described in this paper provides an implementation of such a first step.

Overall, hierarchical approaches are expected to achieve optimal sensitivity at constrained computational resources. We will employ such approaches for deep searches on long duration and high sensitivity data such as those that are now being recorded by the instruments. With 1 year of data at the design sensitivity of the detectors, the improvements that we can expect by means of hierarchical schemes that utilize this type of coherent analysis as one of the steps is of order 10 with respect to what was presented here.

The most constraining 95% confidence level h_0 upper limit from the all-sky search is 6.6×10^{-23} in the band 245.2–246.4 Hz, reflecting the highest sensitivity of the instruments at these frequencies. This is still a factor of $\sim$16 higher than the strongest signal that we expect based on the optimistic (but not unreasonable) assumptions of the statistical argument presented in Sec. III B 1.

The 95% confidence level upper limits from the Sco X-1 search, assuming a noneccentric orbit, are $h_0 \approx 2 \times 10^{-22}$; the most stringent 95% confidence level upper limits from the Sco X-1 search, assuming a noneccentric orbit, are $h_0 = 1.7 \times 10^{-22}$ in the 464–484 Hz band and $h_0 = 2.2 \times 10^{-22}$ in the 604–624 Hz frequency band. The Sco X-1 results presented here are the first direct gravitational wave upper limits placed on the system.

Coherent all-sky searches for continuous signals from isolated stars have been performed in the past, but over much smaller parameter space. In [18] an all-sky 0.76 Hz band search was performed around 921.38 Hz, including spin-down parameters in the range -2.36×10^{-8} Hz s^{-1} to $+2.36 \times 10^{-8}$ Hz s^{-1}. Three data sets, each 48-hours long, were coherently analyzed and a 90% confidence level upper limit was placed at the level of 1.0×10^{-22} based on the cleanest of the data sets.

In the context of a hierarchical search aimed at detecting a signal, data cleaning procedures to remove noise artifacts are likely to be employed. This is in contrast to the approach used in the analysis reported in this paper where no cleaning at all was considered. Here we wanted to investigate the effect of noisy data segments with a variety of artifacts on the output of a search pipeline; we have purposely kept the ''bad'' data in the presentation of the results for illustration purposes. This is relevant to future searches because the upfront ''cleaning'' of known noise artifacts does not guarantee that longer observation times will not uncover unknown periodicities. These must then be either identified as of instrumental origin, as done here, or followed up.

Longer observation times mean a higher resolution in parameter space and higher computational costs. It is thus important to lay the template banks in a way that takes advantage of the correlations in parameter space. The Sco X-1 search presented here already does this. In future work the metric approach will be used also for template placement for the searches for signals from isolated sources. However, longer observation times and targeting different LMXB systems will require more sophisticated template placement strategies than the one presented here for Sco X-1. For example the parameter space may grow to include the orbital period, the eccentricity and the spin period derivatives. For any specific source the number of search parameters will be defined by the precision to which these source parameters have been measured via electromagnetic observations.

Redundant template grids produce redundant events. This increases the false alarm rate and in practice reduces the sensitivity of the search; in fact the threshold on signal-to-noise ratio that defines the candidates to follow up depends by how many follow-ups one can afford with given computational resources. It is therefore important for future work to develop techniques to recognize non-independent candidates in parameter space, rank them, and keep the information on only the most significant. In this paper we have made the first moves in this direction with the algorithm that identifies as a single candidate values of the $\mathcal{F}$ statistic which are ''near'' to each other in search frequency. The concept must be generalized to the multi-dimensional space of the search parameters, and ultimately connected to the global properties of the detection statistic over the parameter space.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory and the Particle Physics and Astronomy Research Council of the United Kingdom, the Max-Planck-Society, and the State of Niedersachsen/Germany for support of the construction and operation of the GEO600 detector. The authors also gratefully acknowledge the support of the research by these agencies and by the Australian Research Council, the Natural Sciences and Engineering Research Council of Canada, the Council of Scientific and Industrial Research of India, the Department of Science and Technology of India, the Spanish Ministerio de Educacion y Ciencia, The National Aeronautics and Space Administration, the John Simon Guggenheim Foundation, the Alexander von Humboldt Foundation, the Leverhulme Trust, the David and Lucile Packard Foundation, the Research Corporation, and the Alfred P. Sloan Foundation. This document has been assigned LIGO Laboratory Document No. LIGO-P050008.

APPENDIX A: SCORPIUS X-1

In this appendix we summarize the observational results that constrain the Sco X-1 parameters and provide justification for the choices made regarding the search area, see IV B 2. We refer to Table I for a summary of the parameters of the source and the relevant errors.

In the analysis we assume the position of the source to be known and we demodulate the data stream for that particular sky location. In fact, the most accurate determination of the Sco X-1 sky position comes from Very Long Baseline Array observations [50,51] and is reported in Table I. The overall error on the source location is $\sim$0.5 arcsec, which is significantly smaller than the $\sim$100 arcsec sky resolution associated with a 6-hour GW search. Hence we assume the position of Sco X-1 (i.e. the barycenter of the binary system) to be exactly known and we point (in software) at that region of the sky.

We turn now to the case of the orbital parameters. In the case of a circular binary ($e = 0$), three are the parameters that describe the orbit and describe the evolution of the Doppler phase shift introduced by the binary motion: the orbital period (P), the projection of the semimajor axis of the orbit $a_{\rm p}$, and what we have defined as the *orbital phase reference time* $\bar{T}$, the time at which the star crosses the ascending node as measured by an observer at the SSB. This is equivalent to setting the argument of periapse (the angle between the ascending node and the direction of periapsis) to zero.

P is by far the most accurately determined parameter [52]; for a coherent search over 6 hours it can be considered known because the loss of signal-to-noise ratio (SNR) introduced by matching two templates with any value of P in the range of Table I is negligible. P becomes a search parameter, requiring multiple filters, only for coherent integration times $\gtrsim 10^6$ s. Much larger uncertainties characterize $a_{\rm p}$ and $\bar{T}$. The large uncertainty on $a_{\rm p}$ is primarily due to the poor determination of the orbital velocity ($40 \pm 5\ {\rm km\,s^{-1}}$ [53]). The uncertainty on the orbital phase reference time is due to the difficulty in locating the Sco X-1 low-mass companion on the orbit. The search therefore requires a discrete grid of filters in the $(a_{\rm p}, \bar{T})$ space.

For the phase model adopted in this analysis we assume that Sco X-1 is in a circular orbit: this is consistent with astrophysical models of semidetached binary systems and with the best fits of the orbital parameters [42]. However, orbital fits for models with $e \neq 0$ were clearly dominated by the noise introduced by the geometry of the Roche lobe [42]. Over an integration time of $\sim$6 h, the eccentricity needs to be smaller than $\sim 10^{-4}$ in order for the detection statistic $\mathcal{F}$ to be affected less than 1%; for $e \approx 10^{-3}$, losses of the order of 10% are expected and are consistent with the results presented in Sec VI B. Unfortunately, current observations are not able to constrain e to such levels of accuracy.

The last parameter we need to search for is the frequency of the gravitational radiation f. The rotation frequency ν of Sco X-1 is inferred from the difference of the frequency of the kHz quasiperiodic oscillations (QPOs). Unfortunately this frequency difference is not constant, and over a 4-day observation [54] has shown a very pronounced drift between 237 ± 5 Hz to 307 ± 5 Hz, where the errors should be interpreted as the 1σ values [55]. This drifting of QPO frequency separation was found to be positively correlated to the inferred mass accretion rate. It is also important to stress that there is a still unresolved controversy as to whether the adopted model that links the rotation frequency to the difference of the frequency of the kHz QPOs is indeed the correct one, and if it is valid for all the observed LMXBs. Moreover, the gravitational wave frequency f is related to ν in a different way, depending on the model that is considered: $f = 2\nu$ if one considers nonaxisymmetric distortions and $f = (4/3)\nu$ if one considers unstable r-modes. It is therefore clear that a search for gravitational waves from Sco X-1 should assume that the frequency is essentially unknown and the whole LIGO sensitivity band (say from ≈ 100 Hz to ≈ 1 kHz) should be considered. Because of the heavy computational burden, such a search requires a different approach (a hierarchical analysis scheme, which is under development). For the analysis presented in this paper, we have decided to confine the search to GWs emitted by nonaxisymmetric distortions ($f = 2\nu$) *and* to constrain the frequency band to the two 20 Hz wide bands (464–484 Hz and 604–624 Hz) that bound the range of the drift of ν.

APPENDIX B: TEMPLATE BANKS

In Sec. V D we have described the filter banks used in the analysis. Here we provide details about the way in which they were generated.

1. Isolated neutron stars

The three-dimensional parameter space is represented by right ascension and declination (α, δ) that identify the source position in the sky and the signal's emission frequency. We construct a filter bank with a 10% maximum mismatch. For simplicity, we employ a uniform grid in frequency, with resolution $\approx 1/T_{\rm span}$, and sky area, as described in Sec. V D 1: the resulting spacing in frequency is $\Delta f_0 = 3.472 \times 10^{-6}$ Hz and an isotropic grid in sky coordinates with equatorial spacing of 0.02 rad, for a total of just under 31 500 filters to cover the entire sky. The number of templates in right ascension α at any given declination δ is proportional to $\cos\delta$.

The grid is chosen based on the maximum expected degradation in the detection statistic due to the mismatch between the actual position of a putative source and the template grid. This effect is measured by Monte Carlo simulations. The simulations consist of series of searches of signals at random locations in the sky with position templates uniformly randomly displaced from the signal's source position by between 0 and half a grid step in both the α and the δ direction. We base the selection of the grid size on the properties of the signals and the simulations are therefore performed in the absence of noise (e.g. see [56]). The grid spacing is chosen in such a way that the expected loss in signal-to-noise ratio due to the signal-template mismatch is a few percent. The results are summarized in Figs. 27 and 28 that refer, respectively, to the L1 and H1 data sets used in the search.

The smaller the maximum mismatch between a signal and a template the more correlated are the filters in the bank. We have estimated the effective number of independent templates from the $2\mathcal{F}$ average loudest event found in single-interferometer searches such as the one described here, in pure Gaussian and stationary noise in 1.2 Hz subbands: 45.7 for L1 and 41.7 for H1. These translate into an effective number of independent templates which is a

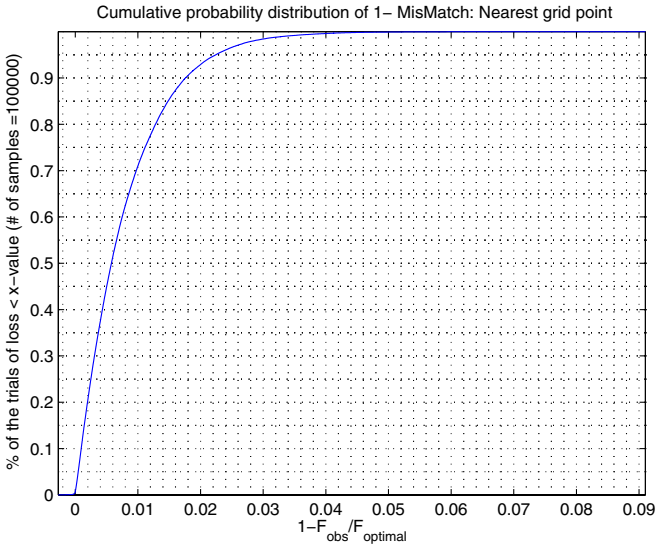

FIG. 27 (color online). Fraction of trials injected in L1 where the ratio $1.0 - \mathcal{F}_{\rm obs}/\mathcal{F}_{\rm optimal}$ is smaller than the value on the x axis. Here in this plot $\mathcal{F}_{\rm obs}$ is the $\mathcal{F}$-statistic at the grid point nearest to the signals' source location. In 99% of the trials the mismatch is smaller than 4%.

factor of ~4 and ~26 smaller for L1 and H1, respectively, than the actual number of templates in the sky grid that we are using for the search. The number of independent templates was estimated as $1/P_0(\mathcal{F}^*)$, where $\mathcal{F}^*$ is the measured loudest value of the detection statistic and P_0 is the false alarm probability defined in Eq. (33). This is consistent with our Monte Carlo simulations (Figs. 27 and 28) where we can see that, with the same grid, 50% of the sky is

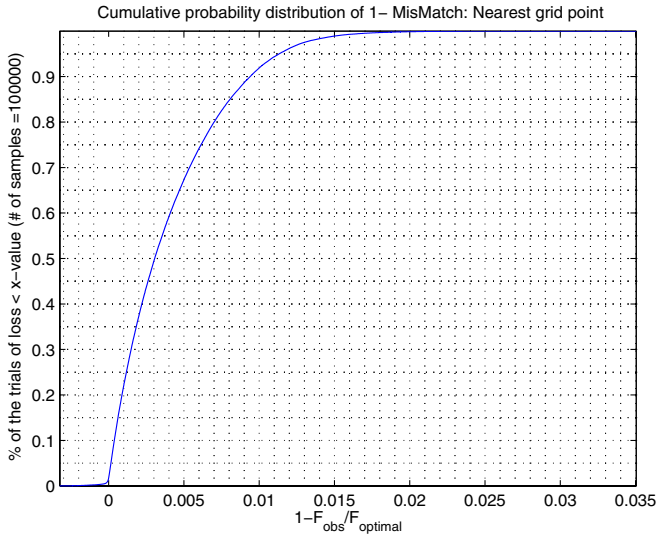

FIG. 28 (color online). Fraction of trials injected in H1 where the ratio $1.0 - \mathcal{F}_{\rm obs}/\mathcal{F}_{\rm optimal}$ is smaller than the value on the x axis. Here in this plot $\mathcal{F}_{\rm obs}$ is the $\mathcal{F}$-statistic at the point nearest to the signals' source location. In 99% of the trials the mismatch is smaller than 2%.

covered in H1 with a mismatch that is always smaller than 0.5% whereas in L1 50% of the sky is covered with a mismatch which is about twice as large. This means that the grid covers H1 data parameter space with more redundancy than it covers L1. The main reason for the difference in coverage is the fact that the spanned observation time of the data set used for the L1 detector is longer than that for H1, and the resolution in sky position is highly dependent on the spanned observation time of the data set.

2. Sco X-1

In the case of Sco X-1 we search over two orbital parameters and the gravitational wave frequency. In order to optimally[7] cover the parameter space, we consider the metric approach introduced in [48] in the context of binary inspirals and applied to pulsar searches in [10,40]. We define the mismatch μ between a signal, described by the parameter vector $\vec{\lambda} = \{f_0, a_{\rm p}, \bar{T}\}$, and a template described by $\vec{\lambda} + \Delta\vec{\lambda}$ as[8]

$$\mu(\vec{\lambda}, \Delta\vec{\lambda}) = 1 - \frac{\mathcal{P}(\vec{\lambda}, \Delta\vec{\lambda})}{\mathcal{P}(\vec{\lambda}, \vec{0})} = g_{\alpha\beta}(\vec{\lambda})\Delta\lambda^\alpha\Delta\lambda^\beta + O(\Delta\lambda^3), \quad \text{(B1)}$$

where the power spectrum $\mathcal{P}$ is given by

$$\mathcal{P}(\vec{\lambda}, \Delta\vec{\lambda}) = \left| \int_0^{T_{\rm span}} e^{-i\Delta\Phi(t;\vec{\lambda},\Delta\vec{\lambda})} dt \right|^2; \quad \text{(B2)}$$

$\Delta\Phi(t; \vec{\lambda}, \Delta\vec{\lambda}) = \Phi(t; \vec{\lambda}) - \Phi(t; \vec{\lambda} + \Delta\vec{\lambda})$ is the difference between the signal and template phase and α, $\beta = 0, 1, 2$ label the search parameters (we follow the convention that the index 0 labels frequency). The metric on the parameter space is given by [10]

$$g_{\alpha\beta} = \langle \partial_\alpha\Delta\Phi\partial_\beta\Delta\Phi\rangle - \langle\partial_\alpha\Delta\Phi\rangle\langle\partial_\beta\Delta\Phi\rangle, \quad \text{(B3)}$$

where $\partial_\alpha\Delta\Phi \equiv \partial\Delta\Phi/\partial\Delta\lambda^\alpha$ and is evaluated at $\Delta\vec{\lambda} = 0$, $\langle\ldots\rangle$ stands for the time average over $T_{\rm span}$, and the gravitational wave phase is defined in Eq. (29). Treating the frequency as a continuous variable (we discuss later in this section the consequences of the fact that the frequency is in practice discrete) and "projecting out" the search frequency dimension of the metric yields a two-dimensional reduced metric only on the orbital parameters,

[7]Although we use the metric approach to lay templates in the parameter space, we use a simple square grid which is non-optimal. Using a hexagonal grid would reduce the number of templates by ~30%.

[8]Note that we use the power spectrum $\mathcal{P}$ to define the mismatch μ and therefore the metric $g_{\alpha\beta}$, but we use the $\mathcal{F}$-statistic in the actual search. As the $\mathcal{F}$-statistic is an optimally weighted sum of squared values of optimally weighted and summed PFTs, we would expect that the template bank is also as effective used with the $\mathcal{F}$-statistic as for the power spectrum. This was tested with extensive software signal injections.

$$\gamma_{jk} = g_{jk} - \frac{g_{0j}g_{k0}}{g_{00}}, \qquad j,k = 1,2. \tag{B4}$$

By using the metric γ_{jk}, we take advantage of the correlations between the frequency and the orbital parameters—so that a mismatch in orbital parameters can be compensated by a mismatch in frequency—and we therefore reduce the number of orbital templates required to cover the parameter space. In the actual analysis, we carry out a coordinate transformation from a_{p} and $\bar{T}$ to two ''search coordinates'' in order to obtain constant spacing and orientation of the filters over the whole parameter space, which simplifies the numerical implementation of the grid.

The frequency, however, is sampled discretely and therefore cannot compensate exactly a mismatch in orbital parameters: it produces a mismatch μ_f only in f between a template and a signal. In order to choose the appropriate frequency spacing Δf we consider the g_{00} component of the metric $g_{\alpha\beta}$—we treat de facto f as a one-dimensional uncorrelated dimension in the parameter space—and obtain

$$\Delta f \approx \frac{2\sqrt{3\mu_f}}{\pi T_{\mathrm{span}}}. \tag{B5}$$

The optimal method for dividing up the *total* mismatch in detection statistic is to split them up equally amongst the dimensions of a parameter space. Therefore, to achieve the required maximal overall mismatch of 10%, we use a 6.6% mismatch in orbital templates and a 3.3% mismatch in frequency. This sets the frequency resolution to $\Delta f = 1/(5T_{\mathrm{span}})$. The grid in the two-dimensional space of the orbital templates is then computed using γ_{jk}, Eq. (B4). In practice, the filter spacing is determined primarily by T_{span} and the grid orientation is determined by the location of the source in its orbit during the observation, see Fig. 10;

As the grid density in the two-dimensional parameter space $(a_{rmp}, \bar{T})$ scales as $(f_0^{\max})^2$, where $f_0^{\max}$ is the maximum search frequency, we generated a different orbital template bank for each 1 Hz frequency sub-band, in order to minimize the computational burden. This reduced by $\approx$ 30% the CPU time that would have been required had we used a single template bank with a maximum frequency parameter $f_0^{\max} = 624$ Hz for the whole analysis. In summary, the number of orbital templates used for each 1 Hz sub-band ranges from 3391 (6681) to 3688 (7236) in the 464–484 Hz band and from 5738 (11309) to 6107 (12032) in the 604–624 Hz band for the L1 (H1) analysis.

APPENDIX C: HARDWARE INJECTIONS

Signals can be injected into the instrument via the actuator, by physically moving the mirrors of a Fabry-Perot cavity to mimic a gravitational wave signal. Hardware injections are designed to give an end-to-end validation of the data analysis pipeline, including some, but not all, components of the calibration. Toward the end of the S2 run, two simulated isolated pulsar signals were injected into the data. We denote the two pulsars P1 and P2 and give their parameters in Table II.

The data sets were prepared using the final S2 calibration version [27], and consist of 17 30-minute SFTs for H1 and 14 30-minute SFTs for L1. We performed a targeted search to look for the pulsar signals.

The results are shown in Table II. The agreement between measured and expected 8NS is good. The measured SNRs, however, are systematically somewhat larger than expected. This is probably due to a small systematic error in the calibration. The differences between the expected SNR values shown here, and those quoted in [24], arise primarily from differences in the lengths of observation times used to make the estimate. In [24], a nominal observation time of 12 hours was used. This is the length of time during which the hardware injections were performed. Here we have used the actual science observation time which is shorter, reflecting science quality data and cali-

TABLE II. Parameters of the two hardware injected pulsars. See Eqs. (25) and (26) for the definition of the parameters. T_0 is the reference time for the initial phase, in GPS seconds in the SSB frame.

	P1	P2
f_0 (Hz)	1279.123 457	1288.901 235
$\dot{f}$ (Hz/s)	0	-10^{-8}
α (rad)	5.147 162	2.345 679
δ (rad)	0.376 696	1.234 568
ψ (rad)	0	0
$\cos\iota$	0	0
Φ_0	0	0
T_0 (sec)	733 967 667.126 112 31	733 967 751.522 490 38
h_0	2×10^{-21}	2×10^{-21}
SNR H1 (exp./meas.)	17/18	34/35
SNR L1 (exp./meas.)	20/22	21/22

bration quality flags based on which we discard data as not reliable enough to be included in an astrophysical search.

APPENDIX D: UPPER LIMITS

In Sec. V G we have described the method used to set upper limits in each of the ~ 1 Hz sub-bands in which the whole frequency bands of the two searches were divided. In this appendix we provide more technical details, on the Monte Carlo approach that we used to determine the value of the gravitational amplitude that we report as upper limit, and the errors that affect such an estimate.

We recall that, given the value of the joint significance of the loudest coincident event, the upper limit on the amplitude h_0^C at confidence level $C(h_0)$ is determined by means of Monte Carlo simulations. In these, a set of N signals all with the same amplitude h_0 is injected in the very same data set used in the analysis and searched for with the same analysis pipeline. For each set of injected signals, and a confidence level $C(h_0)$ is assigned according to Eq. (39), which in turns defines the h_0 upper limit value as a function of the confidence level C.

The details of the Monte Carlo are as follows. For every set of injected signals at fixed h_0, the remaining parameters are chosen randomly from within the boundaries of the parameter space. These include the orbital parameters (orbital semimajor axis and orbital phase reference time) for the Sco X-1 search, the sky position for the all-sky search, and the frequency f_0 and the nuisance parameters ψ, $\cos\iota$, and ϕ_0 for both searches. Uniform distributions are used for f_0 in the sub-band, ψ between 0 and 2π, $\cos\iota$ between -1 and 1, and ϕ_0 between 0 and 2π. For the all-sky search the population of injected signals is uniformly distributed on the celestial sphere, that is to say that α is uniformly distributed between 0 and 2π and $\cos\delta$ uniformly distributed between -1 and 1, with δ between $-\pi/2$ and $\pi/2$. For the Sco X-1 search the signal population is uniformly distributed across the two-dimensional orbital parameter space (the parameters of which are given in Table I). The semimajor axis is selected from within the range 1.26 to 1.62 seconds and the orbital phase reference time is selected from within the GPS time range of 731 163 028 and 731 163 626. The sky position and orbital period are held fixed at values corresponding to the center of their respective ranges (as it was done for the search). For each set of injections the orbital eccentricity is held fixed at one of the following discrete values: $e = 0.0, 10^{-4}, 5 \times 10^{-4}, 10^{-3}, 5 \times 10^{-3}$.

A search over the entire parameter space is *not* performed to search for every injection—it is computationally prohibitive. Rather, the detection statistic is computed *at the nearest template* grid point with respect to the injected signal parameters. The nearest template is chosen consistently with the criteria used for laying out the template bank. For the Sco X-1 search the closest grid point is defined by the metric governing the orbital parameter space. For the isolated pulsar search a Euclidean measure is used. In the actual search noise might conspire to produce a higher value of the detection statistic at a template grid point which is *not* the nearest to the actual signal's parameters. This means that our Monte Carlo may slightly underestimate the detection efficiency of the actual search, leading to an over-conservative (thus still correct) upper limit. However, since our template bank has been chosen so that at most a few percent of the detection statistic may not be recovered at the nearest grid point due to signal-template mismatch, we do not expect this effect to be severe. Furthermore, a detection/coincidence scheme based on the global properties of the detection statistic [57] (far from the signal's true parameters) remains to be understood.

A set of injections at fixed h_0 comprises at least 6000 trials for the isolated pulsar search and 5000 trials for the Sco X-1 search.

To determine the number of injections, several sets of 10 000 isolated pulsar injections have been analyzed. The injections were performed at a fixed strain ($h_0 = 1.2 \times 10^{-22}$) in a small band around 409 Hz, with sky locations and nuisance parameters distributed as described above. Figure 29 shows the results of this analysis. We plot the standard deviation on the confidence level as a function of the number of injections for 7 sets of injections and compare it with the expected values. The plots show that above 5000 injections the standard deviation on the confidence level is less than 0.2% and in agreement with the expectations even for small total number of injections.

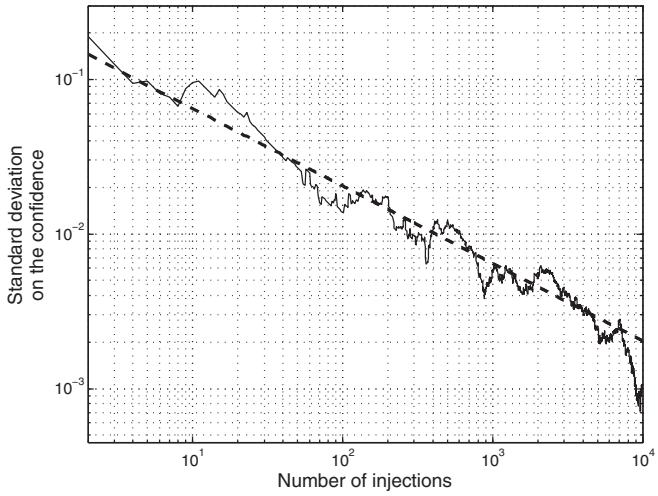

FIG. 29. A total of 7 sets of 10 000 injections was performed in a small band around 409 Hz at a strain $h_0 = 1.2 \times 10^{-22}$. For each set the confidence level was estimated for different numbers of injections. The plot shows the standard deviation based on the 7 estimates as a function of the number of injections. The dashed line shows the expected value of the standard deviation based on the binomial distribution with a single-trial probability of 95.6, which corresponds to the measured mean confidence level at 10 000 injections.

For the isolated pulsar search the following approach has been adopted to estimate the uncertainties related to a frequentist upper limit based on signal injections [19]. The finite sample size of the population of signals that we construct by a Monte Carlo method results in fluctuations in the value of the confidence level C which we measure at fixed gravitational wave amplitude h_0^{injected}. Figure 30 shows a distribution of $(h_0^{\text{injected}}, C)$ values for various sets of injections around the target 95% confidence level value in a reasonably clean sub-band of the data. Close to the target confidence level the relationship $C(h_0^{\text{injected}})$ is well described by a linear relationship. In order to estimate this relation we perform between 6 and 15 sets of injections. Each set is composed of at least 3200 injections and yields a value for $(h_0^{\text{injected}}, C)$. The linear relation is then estimated from these $(h_0^{\text{injected}}, C)$ points with a standard best fit technique. We define $h_0^{95\%}$ as the value of h_0^{injected} yielding $C = 95\%$ according to the fitted linear relation. From the fit we estimate the $\pm 1\sigma$ (h_0, C) curves and from the intercepts of these with $C = 95\%$ the uncertainties on $h_0^{95\%}$, which we expect to be a few percent.

In the Sco X-1 search each set of 5000 injections is divided into 10 subsets, each containing 500 injections. The confidence level $C_i(h_0^{\text{injected}})$ is calculated for each subset of injections where the index $i = 1, \ldots, 10$ labels each subset. Values of $h_{0,i}(C = 95\%)$ are obtained by interpolation between the two values of $C_i(h_0^{\text{injected}})$ closest to 95% within a given subset. The final value of $h_0^{95\%}$ is calculated as the mean of $h_{0,i}(C = 95\%)$ and the uncertainty in this quantity is taken as the standard error in the mean, $\sigma(h_0^{95\%}) = \sqrt{\sigma^2/(10-1)}$, with σ^2 the variance of the $h_{0,i}(C = 95\%)$ sample. This approach has typically yielded uncertainties in the values of the upper limit of $\sim 1\% - 3\%$.

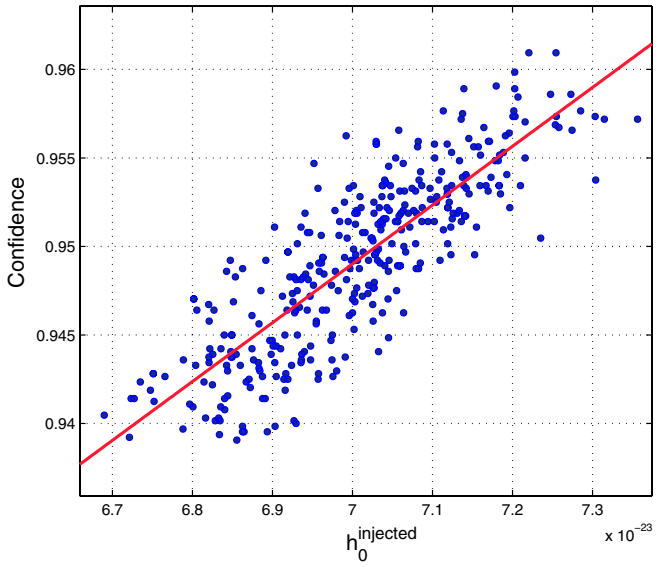

FIG. 30 (color online). Distribution of h_0^{injected} versus confidence level values for sets of 6400 trials. The maximum spread in the confidence level around the target 95% value is about 2% and the maximum spread in h_0 is about 10%. 360 points are plotted in this figure, corresponding to 360 sets, each comprising several thousand injections in the 247.6–248.8 Hz sub-band. The linear fit from these points is $C = 3.3 \times 10^{21} h_0^{\text{injected}} + 0.72$. For $C = 0.95$ the estimated h_0^{injected} is 6.9697×10^{-23}. This is consistent with what is found in the actual analysis with a much smaller (order 10) set of injections: $(6.999 \pm 0.04) \times 10^{-23}$.

The former uncertainties on the upper limit value of h_0 arise only from the Monte Carlo procedure. As we discuss in Sec. V G, an additional and larger error is associated to the instrument calibration.

[1] L. Bildsten, Astrophys. J. **501**, L89 (1998).
[2] G. Ushomirsky, C. Cutler, and L. Bildsten, Mon. Not. R. Astron. Soc. **319**, 902 (2000).
[3] C. Cutler, Phys. Rev. D **66**, 084025 (2002).
[4] A. Melatos and D. J. B. Payne, Astrophys. J. **623**, 1044 (2005).
[5] B. J. Owen, Phys. Rev. Lett. **95**, 211101 (2005).
[6] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson, Phys. Rev. D **58**, 084020 (1998).
[7] N. Andersson, K. D. Kokkotas, and N. Stergioulas, Astrophys. J. **516**, 307 (1999).
[8] D. I. Jones and N. Andersson, Mon. Not. R. Astron. Soc. **331**, 203 (2002).
[9] C. Van Den Broeck, Classical Quantum Gravity **22**, 1825 (2005).
[10] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Phys. Rev. D **57**, 2101 (1998).
[11] P. Brady and T. Creighton, Phys. Rev. D **61**, 082001 (2000).
[12] C. Cutler, I. Gholami, and B. Krishnan, Phys. Rev. D **72**, 042004 (2005).
[13] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **72**, 102004 (2005).
[14] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, Phys. Rev. D **70**, 082001 (2004).
[15] Since February 2005 the LIGO Scientific Collaboration has been operating a public distributed computing project called *Einstein@home*, built using the Berkeley Open Infrastructure for Network Computing (BOINC). Members of the general public can quickly and easily install the software on their Windows, Macintosh, or Linux personal computer. When otherwise idle, their computer downloads data from *Einstein@home*, searches

it for pulsar signals, then uploads information about any candidates. The URL of the project is http://einstein.phys.uwm.edu/.
[16] P. Astone, K. M. Borkowski, P. Jaranowski, and A. Królak, Phys. Rev. D **65**, 042003 (2002).
[17] P. Astone *et al.*, Classical Quantum Gravity **20**, S665 (2003).
[18] P. Astone *et al.*, Classical Quantum Gravity **22**, S1243 (2005).
[19] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **69**, 082004 (2004).
[20] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **69**, 122001 (2004).
[21] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **69**, 102001 (2004).
[22] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **69**, 122004 (2004).
[23] B. Abbott *et al.* (LIGO Scientific Collaboration), Nucl. Instrum. Methods Phys. Res., Sect. A **517**, 154 (2004).
[24] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. Lett. **94**, 181103 (2005).
[25] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **72**, 042002 (2005).
[26] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **72**, 062001 (2005).
[27] G. González, M. Landry, B. O'Reilly, and H. Radkins, LIGO Technical Document T040060-01-D, available at http://admdbsrv.ligo.caltech.edu/dcc/.
[28] C. Kittel, *Introduction to Solid State Physics* (Wiley, Hoboken, 2005), 8th ed..
[29] N. Stergioulas, Living Rev. Relativity **6**, 3 (2003).
[30] R. V. Wagoner, Astrophys. J. **578**, L63 (2002).
[31] N. Andersson, D. I. Jones, and K. D. Kokkotas, Mon. Not. R. Astron. Soc. **337**, 1224 (2002).
[32] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, Astron. J. **129**, 1993 (2005).
[33] K. S. Thorne, in *Three Hundred Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, England, 1987).
[34] F. M. Walter and J. Lattimer, Astrophys. J. **576**, L145 (2002).
[35] A. Tiengo and S. Mereghetti, Astrophys. J. **657**, L101 (2007).
[36] C. Palomba, Mon. Not. R. Astron. Soc. **359**, 1150 (2005).
[37] R. V. Wagoner, Astrophys. J. **278**, 345 (1984).
[38] J. Papaloizou and J. E. Pringle, Mon. Not. R. Astron. Soc. **184**, 501 (1978).
[39] D. Chakrabarty, E. H. Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, M. van der Klis, and C. B. Markwardt, Nature (London) **424**, 42 (2003).
[40] S. V. Dhurandhar and A. Vecchio, Phys. Rev. D **63**, 122001 (2001).
[41] J. H. Taylor and J. M. Weisberg, Astrophys. J. **345**, 434 (1989).
[42] D. Steeghs (private communication).
[43] The searches presented here where mainly carried out on the Merlin cluster at the Albert Einstein Institute, the Tsunami cluster at the University of Birmingham, the Pleiades cluster at Penn State University, and the LDAS cluster at the Hanford LIGO observatory. See http://www.lsc-group.phys.uwm.edu/lscdatagrid/.
[44] P. Jaranowski, A. Królak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998).
[45] S. D. Mohanty, Classical Quantum Gravity **19**, 1513 (2002).
[46] S. D. Mohanty and S. Mukherjee, Classical Quantum Gravity **19**, 1471 (2002).
[47] B. Krishnan, LIGO Technical Document T-040144-00-Z (2004), available at http://admdbsrv.ligo.caltech.edu/dcc/.
[48] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).
[49] See EPAPS Document No. E-PRVDAQ-75-034712 for tables that detail numerically the upper limit results described in the paper. For more information on EPAPS, see http://www.aip.org/pubservs/epaps.html.
[50] The Very Long Baseline Array Telescope, http://www.vlba.nrao.edu/.
[51] C. F. Bradshaw, E. B. Fomalont, and B. J. Geldzahler, Astrophys. J. **512**, L121 (1999).
[52] E. W. Gottlieb, E. L. Wright, and W. Liller, Astrophys. J. **195**, L33 (1975).
[53] D. Steeghs and J. Cesares, Astrophys. J. **568**, 273 (2002).
[54] M. van der Klis, R. A. D. Wijnands, K. Horne, and W. Chen, Astrophys. J. Lett. **481**, L97 (1997).
[55] M. van der Klis, Annu. Rev. Astron. Astrophys. **38**, 717 (2000).
[56] B. Owen and B. S. Sathyaprakash, Phys. Rev. D **60**, 022002 (1999).
[57] R. Prix and Y. Itoh, Classical Quantum Gravity **22**, S1003 (2005).