Rapid localization of compact binaries with Advanced LIGO

221st AAS Meeting • 9 Jan 2013 • Long Beach, CA

Leo Singer & Larry Price

lsinger@caltech.edu • larryp@caltech.edu

Advanced LIGO and Virgo should secure the first direct detections of gravitational waves. Compact binary mergers, promising gravitational-wave sources and possible progenitors of short gamma-ray bursts, are predicted to have rapidly fading optical counterparts. Optical follow-ups will demand timely and accurate parameter estimation. In the last LIGO—Virgo science run, a triangulation code produced sky maps for pointing telescopes. The improved Bayesian sky map algorithm described here will complement the exhaustive but higher latency Markov-chain Monte Carlo parameter estimation.

Problem statement

Compact binary mergers are the most promising sources for ground-based gravitational-wave detectors. Triangulation on time-of-arrival and amplitude gives a prompt estimate of the sky location of a gravitational-wave event suitable for pointing telescopes to search for faint, fast transients [1]. Given adequate time and computing resources, powerful Monte Carlo codes further constrain sky location, masses, and spins straight from the gravitational-wave data [2]. We envision both approaches as two tiers of one follow-up effort with time-evolving precision.

To fill the first tier, we are developing a rapid Bayesian sky localization code that uses the time and amplitude estimated by low-latency detection pipelines, but is faster and more accurate than the triangulation code that was developed for the last joint LIGO—Virgo science run.

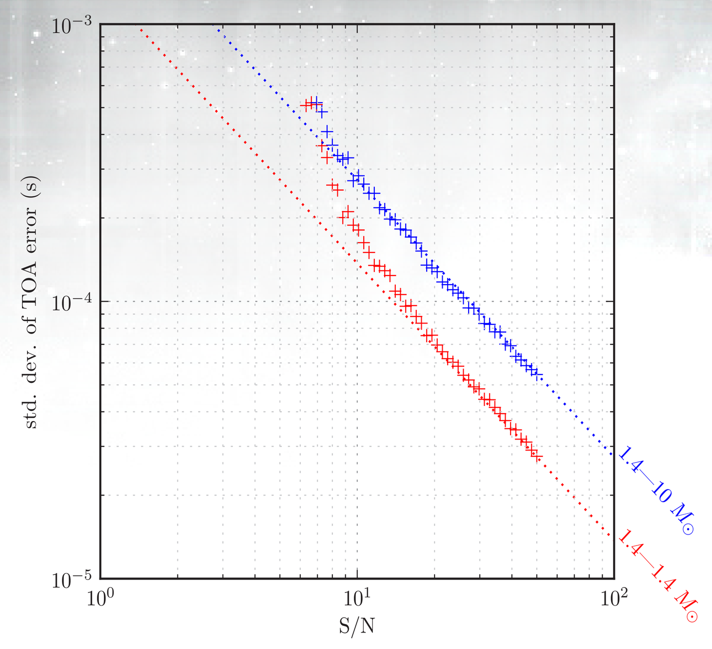

Figure 1—As a function of SNR, timing accuracy of a matched filter (crosses) compared with the Cramér-Rao bound (dotted lines).

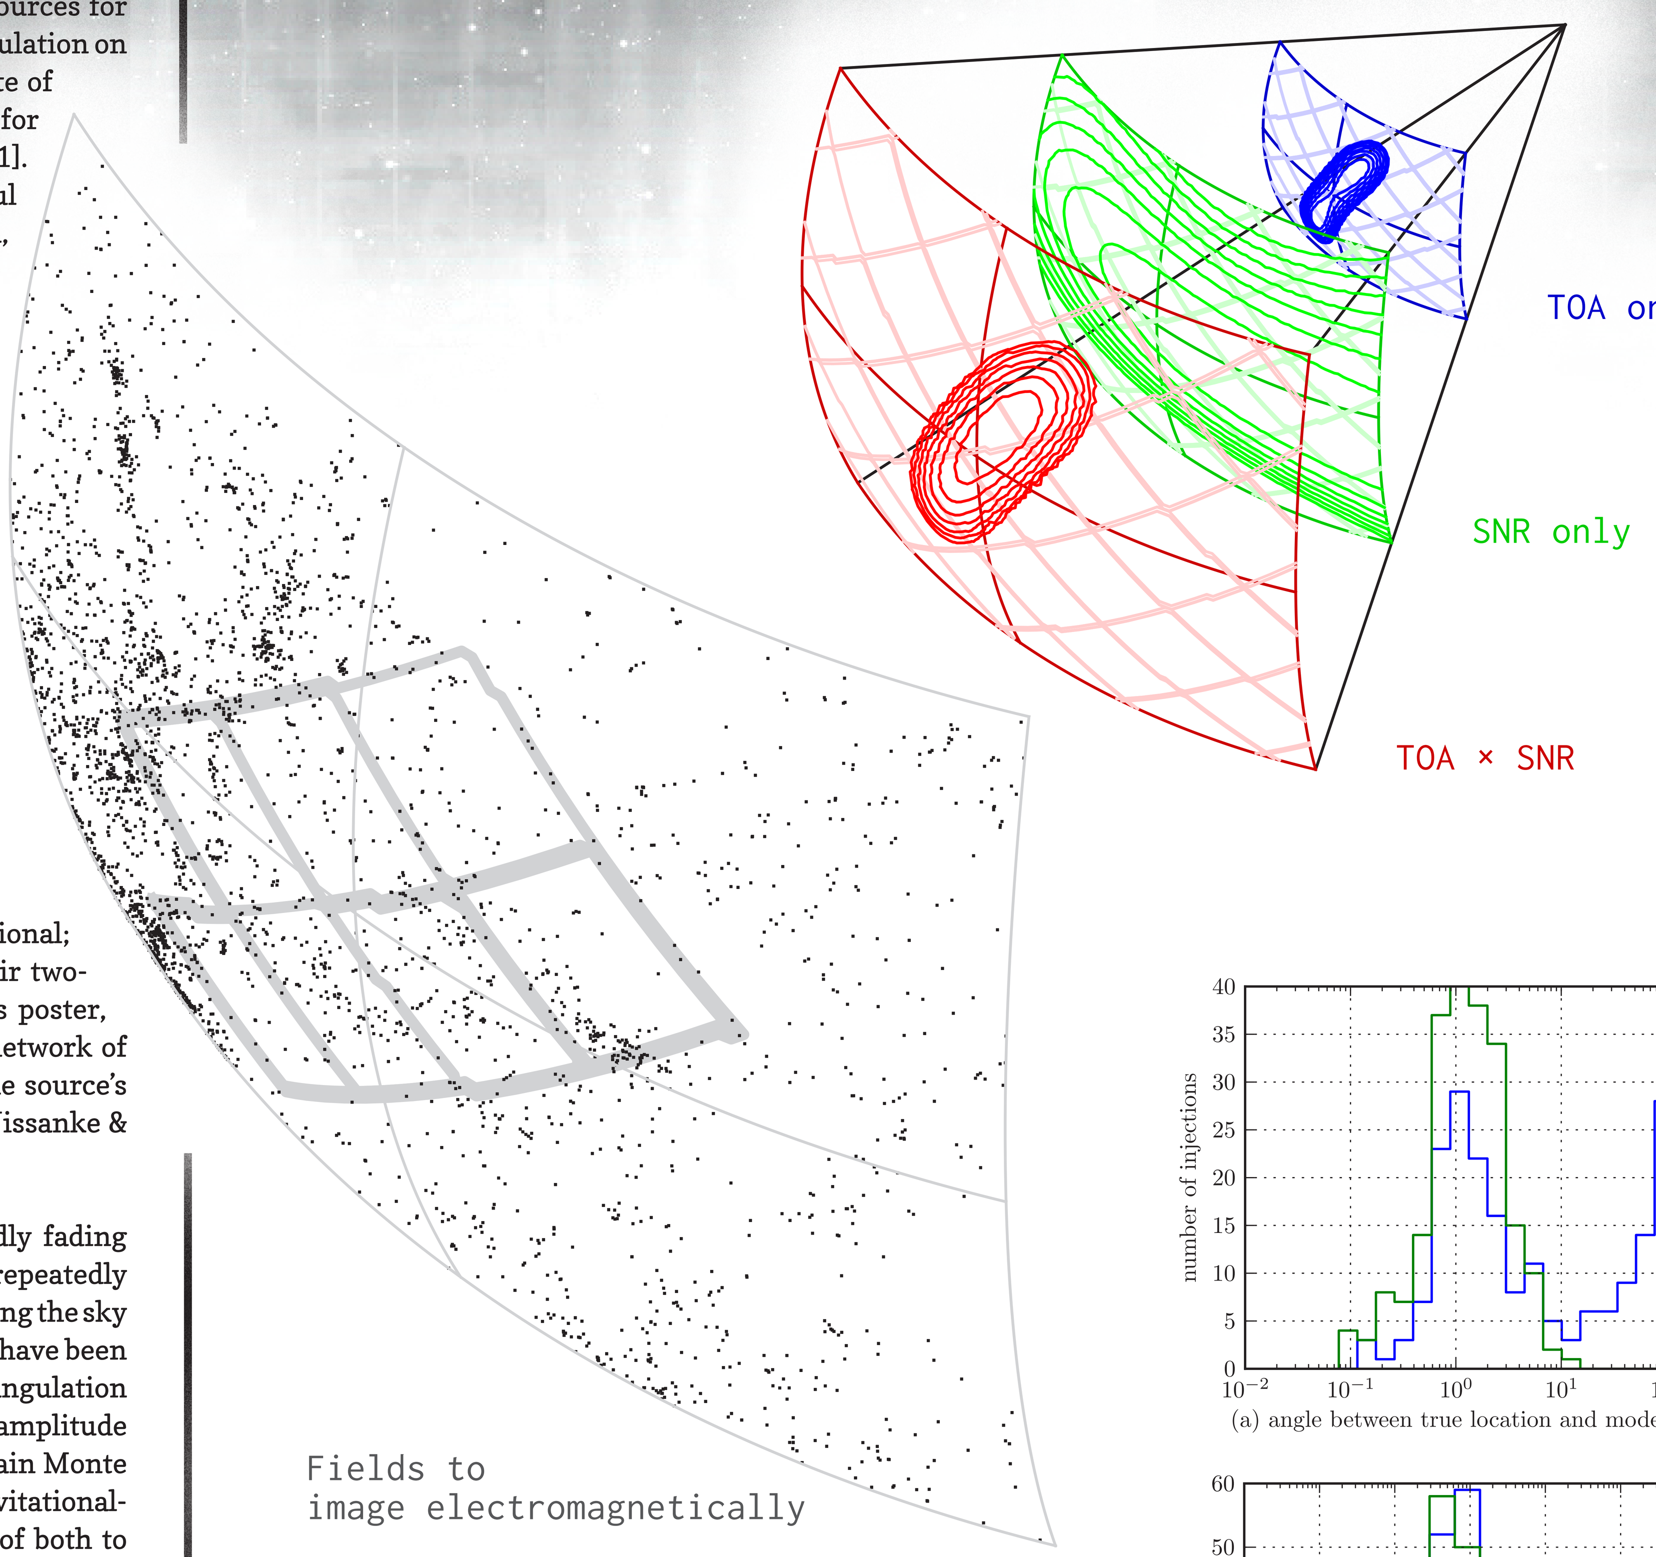

Astrophysical motivation

As a compact binary system loses energy to gravitational radiation, its orbit decays, ultimately driving the components of the binary to merge. In its final moments, a neutron star—neutron star, neutron star—black hole, or black hole—black hole binary, becomes an efficient source of gravitational radiation in frequencies that are ideal for ground-based gravitational wave observatories such as LIGO and Virgo, from tens to hundreds of Hz. In their final 'Advanced' configuration, these facilities may realistically detect tens of mergers per year [3].

If one of the components of the system is a neutron star, it may be tidally shredded before merging, providing fuel for an electromagnetic counterpart [4]. The short-lived accretion flow may power a collimated relativistic jet, resulting in a short gamma-ray burst and an X-ray afterglow if it is aligned with the line of sight. An optical afterglow may follow minutes to hours later as the jet plows into the interstellar medium. As surrounding neutron-rich ejecta decay radioactively, an optical 'kilonova' may be visible after ~1 day. Finally, as the ejecta plow into the interstellar medium, a faint radio afterglow will occur on a timescale of years.

Limitations of instruments

Gravitational wave detectors are omnidirectional; their only directional sensitivity comes from their two-polarization antenna patterns [refer to V. Raymond's poster, 346.06]. An observation of a compact binary with a network of gravitational wave detectors will seldom constrain the source's sky location to better than ~10-100 deg^2 [5, and see Nissanke & Kasliwal's poster, 346.03].

There are many challenges in locating a faint, rapidly fading optical transient in such a large area. Tiling the region repeatedly with multiple optical observations will require estimating the sky location accurately and rapidly. Two broad techniques have been used so far for gravitational-wave sky localization: triangulation proceeding from matched-filter time-of-arrival and amplitude in each detector, and exhaustive Bayesian Markov-chain Monte Carlo parameter estimation proceeding from the gravitational-wave signal itself. We propose to combine elements of both to form a fast, Bayesian sky localization algorithm that takes the matched-filter times and amplitudes as input.

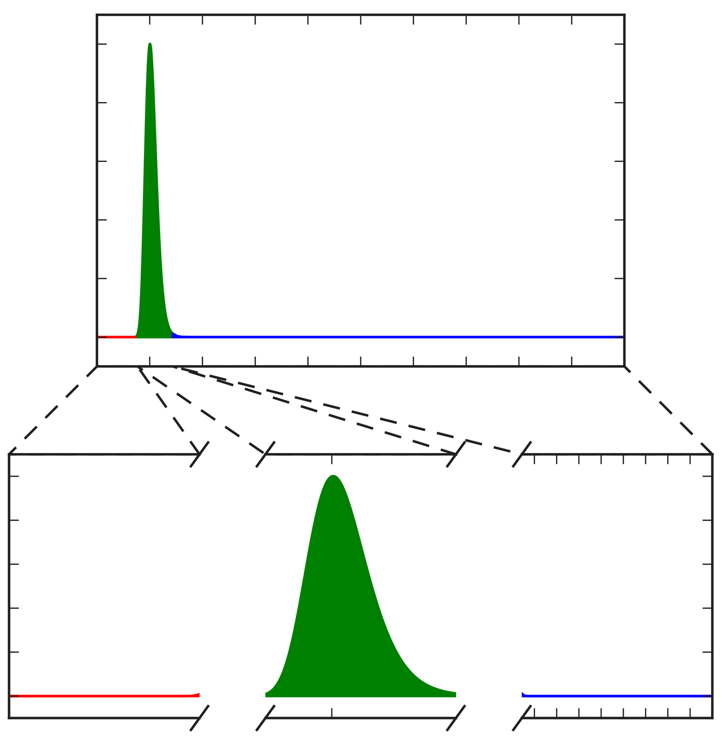

Figure 2—Illustration of adaptive quadrature scheme. The top panel shows the radial integrand. The limits of integration are initially divided at the points where the likelihood attains 1% of its maximum value. The integrator then adaptively subdivides those intervals to attain a given accuracy criterion. The initial division makes it so that the integrator quickly finds the most significant region.

Algorithm

Matched filters provide point estimates of the time, amplitude, and phase of the signal at each detector. The Cramér-Rao bound has been used extensively to model the parameter estimation accuracy of gravitational-wave observations [6, 7]. For real amplitude ρ, phase γ, and time delay τ, and denoting parameter estimates with hats and estimate errors with tildes ($\tilde{\tau} = \tau - \hat{\tau}$, etc.), the Cramér-Rao bound is

$$\mathrm{E}\left[\begin{pmatrix}\tilde{\rho}\\ \tilde{\gamma}\\ \tilde{\tau}\end{pmatrix}\begin{pmatrix}\tilde{\rho}\\ \tilde{\gamma}\\ \tilde{\tau}\end{pmatrix}^{\mathsf{T}}\right] \geq \begin{pmatrix}1 & 0 & 0\\ 0 & \rho^2 & -\rho^2\overline{\omega^1}\\ 0 & -\rho^2\overline{\omega^1} & \rho^2\overline{\omega^2}\end{pmatrix}^{-1}$$

where $\overline{\omega^k} \equiv \left(\frac{2}{\pi}\int_0^\infty \frac{|u(\omega)|^2}{S_n(\omega)}\omega^k \mathrm{d}\omega\right)\left(\frac{2}{\pi}\int_0^\infty \frac{|u(\omega)|^2}{S_n(\omega)}\mathrm{d}\omega\right)^{-1}.$

We infer that amplitude errors are uncorrelated with time and phase (though time and phase errors are correlated with each other). Consequently, we factor the likelihood into a time-of-arrival (TOA) part and a signal-to-noise-ratio (SNR) part, $\mathcal{L} = \mathcal{L}_{\mathrm{TOA}} \times \mathcal{L}_{\mathrm{SNR}}$. We treat the TOA part as Gaussian with mean depending on sky location $\mathbf{n}$ and absolute arrival time $t_\oplus$:

$$\mathcal{L}_{\mathrm{TOA}} \propto \exp\left[-\frac{1}{2}\sum_i \frac{\left(t_i(\mathbf{n}, t_\oplus) - \hat{t}_i\right)^2}{\sigma_{t_i}^2}\right]$$

and the SNR part as Gaussian with mean depending on location, distance, inclination, polarization angle, and coalescence phase:

$$\mathcal{L}_{\mathrm{SNR}} \propto \exp\left[-\frac{1}{2}\sum_i \left(\rho_i(\mathbf{n}, D_{\mathrm{L}}, \iota, \psi, \phi_c) - \hat{\rho}_i\right)^2\right]$$

With a uniform prior in $t_\oplus$, $\cos\iota$, ψ, ϕ_c, and either $\log D_{\mathrm{L}}$ or volume, we marginalize over the following nuisance parameters:

$t_\oplus$, ϕ_c — analytically
ι, ψ — Newton–Cotes
D_{L} — adaptive Gaussian quadrature

Hierarchical refinement

We first evaluate the TOA-only sky map. Ranking its pixels by decreasing probability, we compute the costly SNR term for just those pixels that comprise 0.9999 of the TOA posterior.

Adaptive integration

So that the distance integral converges rapidly, we subdivide the integration limits into up to four intervals that enclose the sides of the maximum-likelihood peak and the tails of the likelihood as illustrated in Figure 2. These intervals are initial subdivisions for an adaptive integrator [`gsl_integrate_qagp`, 8].

The integrator automatically spends more time near the maximum-likelihood peak when the posterior is likelihood-dominated or more time in the tails when prior-dominated.

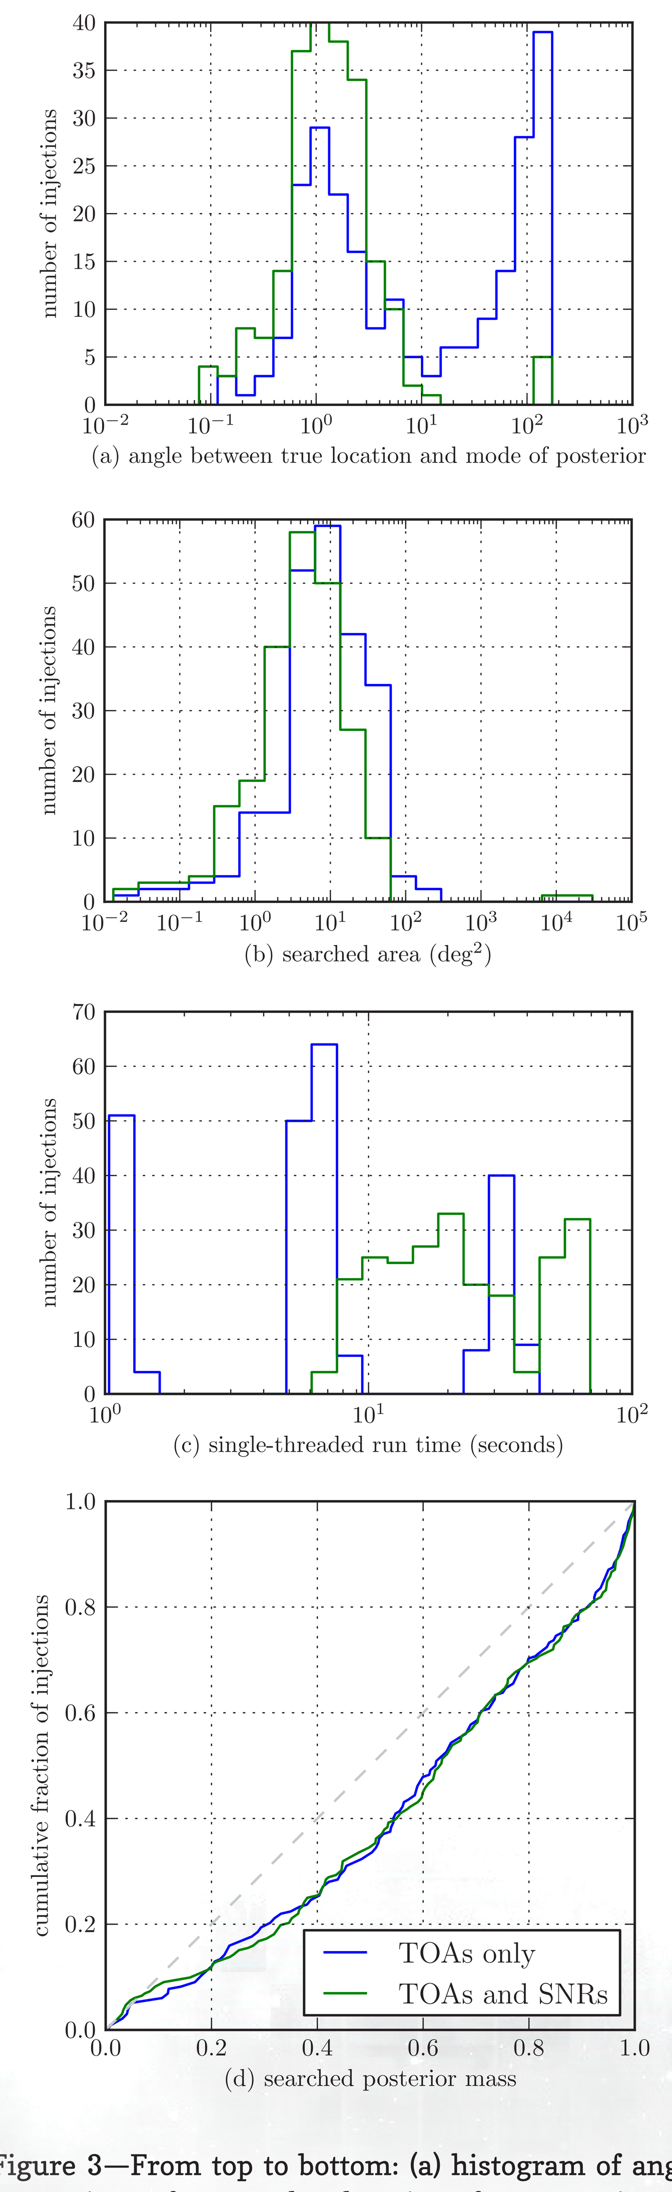

Figure 3—From top to bottom: (a) histogram of angular separation of true sky location from maximum a posteriori estimate, (b) histogram of area enclosed by smallest posterior probability contour containing true sky location, (c) histogram of single-threaded run time, and (d) histogram of smallest percentile of sky map containing true sky location.

Performance

We have measured the performance of our new sky localization code on a population of 244 canonical 1.4—1.4 $M_\odot$ simulated neutron star—neutron star signals in Gaussian noise, logarithmically distributed in distance from 100 to 300 Mpc, and restricted to SNR ≥ 8 in each of three detectors. For each simulation, we generated a sky map using time-of-arrival only and then using time-of-arrival and amplitude.

In Figure 3 at right, we summarize the performance on these simulations. Figure 3(a) is a histogram of the angle between the true sky location and the maximum a posteriori estimate. Figure 3(b) shows the area in square degrees of the lowest confidence region that contains the true location, representative of the area on the sky that would have to be imaged in order to reach the candidate. Figure 3(c) shows the run time using a single thread. Figure 3(d) shows the cumulative fraction of events that are localized within a given confidence level.

From Figures 3(a) and 3(b), we see that adding the amplitude information reduces the searched area by a factor of two by breaking the degeneracy that occurs with time-of-arrival-only triangulation with fewer than four detectors. Figure 3(c) shows that the sky localization process takes at most several tens of seconds on a single core (speedup is 1/number of cores). Figure 3(d) demonstrates that the algorithm nearly satisfies the necessary condition that the sky localization code produces self-consistent probability contours such that p percent of sources are found in the pth percentile of their respective sky maps. Since the data dips below the diagonal, our probability contours are slightly overconfident. Further investigation is needed.

Conclusion

We have combined an ideal matched-filter detection pipeline with a semi-analytic Bayesian parameter estimation scheme in order to rapidly infer the sky locations of simulated gravitational wave signals from binary neutron star mergers.

In the future, we will test our algorithm in real LIGO/Virgo noise and then incorporate it into existing parameter estimation codes. Applications beyond rapid localization include providing jump proposals for the full Monte Carlo parameter estimation.

References

[1] Abadie, J. et al. 2012. First low-latency LIGO+Virgo search for binary inspirals and their electromagnetic counterparts. A&A, 541, A155.

[2] Abadie, J. et al. 2012. Parameter estimation for compact binary coalescence events with first-generation ground-based gravitational-wave interferometers. LIGO Document P1200021-v6.

[3] Abadie, J. et al. 2010. Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. CQG, 27, 173001.

[4] Metzger, B. & Berger, E. 2012. What is the most promising electromagnetic counterpart of a neutron star binary merger? ApJ, 746, 48.

[5] Nissanke, S. et al. 2012. Identifying elusive electromagnetic counterparts to gravitational wave mergers: an end-to-end simulation. ApJ, in press.

[6] Cutler, C. & Flanagan, É. 1994. Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters? PRD, 49, 2658.

[7] Fairhurst, S. 2009. Triangulation of gravitational wave sources with a network of detectors. New J. Phys., 11, 123006.

[8] <http://www.gnu.org/software/gsl/manual/html_node/QAGP-adaptive-integration-with-known-singular-points.html>

LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation (NSF) and operates under cooperative agreement PHY-0107417. This research is supported by NSF through a Graduate Research Fellowship to LS. This is LIGO Document G1200926-v4.