Pulsar timing arrays as imaging gravitational wave telescopes: angular resolution and source (de)confusion

Latham Boyle[1] and Ue-Li Pen[2]
[1] *Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada*
[2] *Canadian Institute for Theoretical Astrophysics, Toronto, Ontario, Canada*
(Dated: October 2010)

Pulsar timing arrays (PTAs) will be sensitive to a finite number of gravitational wave (GW) "point" sources (e.g. supermassive black hole binaries). N quiet pulsars with accurately known distances d_{pulsar} can characterize up to $2N/7$ distant chirping sources per frequency bin $\Delta f_{gw} = 1/T$, and localize them with "diffraction limited" precision $\delta\theta \gtrsim (1/\mathrm{SNR})(\lambda_{gw}/d_{pulsar})$. Even if the pulsar distances are *poorly* known, a PTA with F frequency bins can still characterize up to $(2N/7)(1-\frac{1}{2F})$ sources per bin, and the quasi-singular pattern of timing residuals in the vicinity of a GW source still allows the source to be localized quasi-topologically within roughly the smallest quadrilateral of quiet pulsars that encircles it on the sky, down to a limiting resolution $\delta\theta \gtrsim (1/\mathrm{SNR})\sqrt{\lambda_{gw}/d_{pulsar}}$. PTAs may be unconfused, even at the lowest frequencies, with matched filtering always appropriate.

Our Local Group of galaxies is sprinkled with millisecond pulsars – natural clocks of extraordinary stability. Gravitational waves (GWs) passing through the Milky Way, after being generated *e.g.* by the inspiral of two supermassive black holes in a distant galaxy, generate fluctuations in the time of arrival (TOA) of the pulses at the Earth [1, 2]. In the future, we are likely to detect such GWs via their coherent imprint on the TOA fluctuations from a collection of pulsars distributed on the sky: a "pulsar timing array" (PTA). Much research has focused on using PTAs to study stochastic GW backgrounds ([3–5] and references therein). Recently, various authors have begun to study the ability of PTAs to detect and characterize individual GW point sources [6–13].

Continuing in this direction, this paper is concerned with conceptually clarifying the theoretical behavior and capabilities of PTAs as GW point source telescopes. We address two related issues. (i) A PTA may be sensitive to so many GW sources that it becomes "confused" – *i.e.* unable to disentangle and individually characterize the sources. When does a PTA become "confusion limited" rather than sensitivity limited? How many GW sources is it capable of individually characterizing? (ii) When a set of GW point sources can be individually characterized, how well can their angular positions be determined?

Regarding issue (i) we will see that PTAs with many pulsars can characterize many GW sources per frequency bin; the traditional rule of thumb that a gravitational wave detector becomes confused when there is more than about one source per bin is too pessimistic for PTAs. Regarding issue (ii) we must distinguish pulsars whose distances are known *accurately* or *poorly* relative to $\lambda_{gw}/(1+\cos\theta)$, where λ_{gw} is the gravitational wavelength and θ is the angle between pulsar and source. Pulsars with accurately known distances can angularly localize a GW source very precisely; each such pulsar acts like a single baseline of a diffraction-limited radio interferometer array – with the radio wavelength replaced by the gravitational wavelength, and the length of the radio baseline replaced by the distance from the pulsar to the Earth! The contribution from pulsars with poorly known distances is more interesting: due to a quasi-singularity in the pattern of timing residuals near the location of the GW source, the source can still be localized surprisingly well, for reasons that have less to do with diffraction, and more to do with topology!

Basic Formalism. We will label the 3 spatial directions with the latin indices $\{i, j, k, l, m = 1, 2, 3\}$, raised and lowered with δ^{ij} and δ_{ij}. The N pulsars in the network are labelled by the greek indices $\{\alpha, \beta = 1, \ldots, N\}$, raised and lowered with $\delta^{\alpha\beta}$ and $\delta_{\alpha\beta}$. We follow the Einstein summation convention: repeated indices (one upper, one lower) are summed.

A gravitational wave on Minkowski space is described in transverse-traceless (TT) gauge [14] by the line element $ds^2 = -dt^2 + [\delta_{ij} + 2h_{ij}]dx^i dx^j$. In this gauge, the $\vec{x}$ = constant worldlines are timelike geodesics; along such worldlines, the proper time τ is the coordinate time t. To avoid notational clutter, let us start with just a single gravitational plane wave travelling in the $\hat{n}$ direction:

$$h_{ij}(t,\vec{x}) = \int_{-\infty}^{\infty} df \tilde{h}_{ij}(f) \mathrm{e}^{2\pi i f(\hat{n}\cdot\vec{x}-t)}; \qquad (1)$$

it is straightforward to extend the following analysis to a sum of $m = 1, \ldots, M$ plane waves, each travelling in a different direction $\hat{n}_m$; this extension is discussed below. Throughout this paper, we use "dot product" notation to mean contraction with the unperturbed 3-metric δ_{ij}: $\vec{a}\cdot\vec{b} \equiv \delta_{ij}a^i b^j$; and hats denote unit 3-vectors: $\hat{a}\cdot\hat{a} = 1$.

If an electromagnetic flash is emitted from position $\vec{x}_i$ at time t_i, what is its arrival time t at position $\vec{x}_f$? If we define $\vec{x}_{fi} \equiv \vec{x}_f - \vec{x}_i = x_{fi}\hat{x}_{fi}$ then, at *zeroth* order (*i.e.* in the absence of gravitational waves) the answer is $t_0 = t_i + x_{fi}$. Solving the geodesic equation to first order in h_{ij} yields the perturbed result $t = t_0 + \delta t$ where:

$$\delta t = \int df \frac{i\tilde{h}_{ij}(f)\hat{x}^i_{fi}\hat{x}^j_{fi}[\mathrm{e}^{2\pi i f(\hat{n}\cdot\vec{x}_f - t_0)} - \mathrm{e}^{2\pi i f(\hat{n}\cdot\vec{x}_i - t_i)}]}{2\pi f(1-\hat{n}\cdot\hat{x}_{fi})}. \qquad (2)$$

Now consider an observer at fixed spatial position $\vec{x}=\vec{0}$ receiving signals from $\alpha=1,\ldots,N$ pulsars at spatial positions $\vec{r}_\alpha=r_\alpha\hat{r}_\alpha$. For pulsar α, the TOA fluctuation $\delta t_\alpha(t_0)$, as a function of the unperturbed TOA t_0, is

$$\delta t_\alpha(t_0)=\int_{-\infty}^{\infty}df\delta\tilde{t}_\alpha(f)\mathrm{e}^{-2\pi ift_0} \tag{3}$$

where

$$\delta\tilde{t}_\alpha(f)=\frac{i\tilde{h}_{ij}(f)\hat{r}^i_\alpha\hat{r}^j_\alpha[1-\mathcal{P}_\alpha(f)]}{2\pi f(1+\hat{n}\cdot\hat{r}_\alpha)} \tag{4}$$

and, for later convenience, we have defined the phase

$$\mathcal{P}_\alpha(f)\equiv\mathrm{e}^{2\pi ifr_\alpha(1+\hat{n}\cdot\hat{r}_\alpha)}. \tag{5}$$

The *measured* TOA fluctuations $s_\alpha(t_0)$ from pulsar α are gravitational wave signal $\delta t_\alpha(t_0)$ plus noise $n_\alpha(t_0)$:

$$s_\alpha(t_0)=\delta t_\alpha(t_0)+n_\alpha(t_0). \tag{6}$$

We take the noise to be stationary and gaussian, so it is characterized by its correlation function $C_{\alpha\beta}(T)$ or, equivalently, its spectral density $S_{\alpha\beta}(f)=\tilde{C}_{\alpha\beta}(f)$:

$$C_{\alpha\beta}(T) \;=\; \overline{n_\alpha(t_0+T)n_\beta(t_0)} \tag{7a}$$
$$\delta(f-f')S_{\alpha\beta}(f) \;=\; \overline{\tilde{n}^*_\alpha(f)\tilde{n}_\beta(f')}. \tag{7b}$$

We also take the noise to be uncorrelated between different pulsars: $S_{\alpha\beta}(f)=S_\alpha(f)\delta_{\alpha\beta}$. Let us define the natural noise-weighted inner product:

$$(g^{(1)}|g^{(2)})=\int_{-\infty}^{\infty}df\,\tilde{g}^{(1)}_\alpha(f)^*[S^{-1}(f)]^{\alpha\beta}\tilde{g}^{(2)}_\beta(f). \tag{8}$$

Then matched filtering will detect a given gravitational wave signal with expected signal-to-noise ratio squared (SNR^2) given by

$$\mathrm{SNR}^2=(\delta t|\delta t)=\sum_{\alpha=1}^{N}\mathrm{SNR}^2_\alpha \tag{9a}$$
$$\mathrm{SNR}^2_\alpha=\int_{-\infty}^{\infty}df\frac{|\delta\tilde{t}_\alpha(f)|^2}{S_\alpha(f)} \tag{9b}$$

where $\delta\tilde{t}_\alpha(f)$ is given by (4). When a gravitational wave signal (which depends on various parameters ξ^k) is detected with sufficient SNR, the likelihood function may be approximated as a gaussian $\propto\exp[-(1/2)\xi^k\Gamma_{kl}\xi^l]$ near its peak, and the expected inverse covariance matrix is the Fisher information matrix, given by

$$\Gamma_{kl}=\Big(\frac{\partial t}{\partial\xi^k}\Big|\frac{\partial t}{\partial\xi^l}\Big). \tag{10}$$

We are interested, in particular, in the angular resolution of a PTA. Define an orthonormal triad from $\hat{n}$ and two other unit vectors $\hat{m}_{\bar{\mu}}$ ($\bar{\mu}=1,2$); let $\gamma^{\bar{\mu}}$ be the rotation angle around $\hat{m}_{\bar{\mu}}$. The 2×2 angular part of Γ_{kl} is

$$\Gamma_{\bar{\mu}\bar{\nu}}=\Big(\frac{\partial[\delta t]}{\partial\gamma^{\bar{\mu}}}\Big|\frac{\partial[\delta t]}{\partial\gamma^{\bar{\nu}}}\Big)=\sum_{\alpha=1}^{N}\Gamma^\alpha_{\bar{\mu}\bar{\nu}} \tag{11a}$$
$$\Gamma^\alpha_{\bar{\mu}\bar{\nu}}=\int_{-\infty}^{\infty}df\frac{\partial[\delta\tilde{t}_\alpha(f)]}{\partial\gamma^{\bar{\mu}}}^*\frac{1}{S_\alpha(f)}\frac{\partial[\delta\tilde{t}_\alpha(f)]}{\partial\gamma^{\bar{\nu}}}. \tag{11b}$$

To evaluate these angular derivatives, we act with the infinitessimal rotation $R_{ij}\approx\delta_{ij}-\epsilon_{ijk}\hat{m}^k_{\bar{\mu}}\gamma^{\bar{\mu}}$ on the gravitational wave field, but *not* on the pulsar positions: *e.g.* $\partial(1+\hat{n}\cdot\hat{r}_\alpha)/\partial\gamma^{\bar{\mu}}=\epsilon_{ijk}\hat{n}^i\hat{r}^j_\alpha\hat{m}^k_{\bar{\mu}}$ and $\partial[\tilde{h}_{ij}(f)\hat{r}^i_\alpha\hat{r}^j_\alpha]/\partial\gamma^{\bar{\mu}}=2\tilde{h}_{il}(f)\epsilon_{jk}{}^{l}\hat{r}^i_\alpha\hat{r}^j_\alpha\hat{m}^k_{\bar{\mu}}$. In this way, we find

$$\frac{\partial[\delta\tilde{t}_\alpha(f)]}{\partial\gamma^{\bar{\mu}}}=\frac{i[\mathcal{A}(f)+\mathcal{B}(f)]}{2\pi f} \tag{12}$$

where $\mathcal{A}$ comes from differentiating the phase $\mathcal{P}_\alpha$ in (4), and $\mathcal{B}$ comes from differentiating everything else:

$$\mathcal{A}\equiv2\pi ifr_\alpha\frac{\tilde{h}_{ij}\hat{r}^i_\alpha\hat{r}^j_\alpha\hat{r}^k_\alpha\hat{n}^l\hat{m}^m_{\bar{\mu}}\epsilon_{klm}}{(1+\hat{n}\cdot\hat{r}_\alpha)}\mathcal{P}_\alpha \tag{13a}$$
$$\mathcal{B}\equiv\frac{\tilde{h}_{ij}\hat{r}^i_\alpha\hat{r}^k_\alpha\hat{m}^l_{\bar{\mu}}}{1+\hat{n}\cdot\hat{r}_\alpha}\Big[2\epsilon^j{}_{kl}-\frac{\hat{r}^j_\alpha\hat{n}^m\epsilon_{klm}}{1+\hat{n}\cdot\hat{r}_\alpha}\Big][1-\mathcal{P}_\alpha]. \tag{13b}$$

In understanding the meaning of these equations, we should distinguish two cases: (i) pulsars whose distances r_α are known *accurately* relative to $\lambda_{gw}/(1+\hat{n}\cdot\hat{r}_\alpha)$, so $\mathcal{P}_\alpha$ is known; and (ii) pulsars whose distances r_α are known *poorly* relative to $\lambda_{gw}/(1+\hat{n}\cdot\hat{r}_\alpha)$, so $\mathcal{P}_\alpha$ is essentially a random phase. We consider these two cases in turn.

Pulsars with accurately known distances. First consider a monochromatic gravitational plane wave of frequency $f_{gw}=c/\lambda_{gw}$, and a pulsar whose distance r_α is known accurately relative to $\lambda_{gw}/(1+\hat{n}\cdot\hat{r}_\alpha)$. Then, if $2\pi r_\alpha/\lambda_{gw}\gg1$, the $\mathcal{A}$ term dominates the $\mathcal{B}$ term in Eq. (12) and we have

$$\frac{\Gamma^\alpha_{\bar{\mu}\bar{\nu}}}{\mathrm{SNR}^2_\alpha}\approx\Big(\frac{2\pi r_\alpha}{\lambda_{gw}}\Big)^2\frac{(\hat{r}^i_\alpha\hat{n}^j\hat{m}^k_{\bar{\mu}}\epsilon_{ijk})(\hat{r}^{i'}_\alpha\hat{n}^{j'}\hat{m}^{k'}_{\bar{\nu}}\epsilon_{i'j'k'})}{|1-\mathcal{P}_\alpha(f_{gw})|^2}. \tag{14}$$

Since the second fraction on the right-hand side of this equation is typically $\mathcal{O}(1)$, this says that when a pulsar at a well known distance r_α registers a gravitational wave with signal-to-noise level SNR_α, its contribution to $\Gamma^\alpha_{\bar{\mu}\bar{\nu}}$ is typically $\Gamma^\alpha_{\bar{\mu}\bar{\nu}}\sim(2\pi r_\alpha/\lambda_{gw})^2\mathrm{SNR}^2_\alpha$. In other words, each such pulsar acts is just like one of the baselines of a radio interferometer array; but, in this analogy, the radio waves are replaced by gravitational waves, and the baselines are of galactic length scales and extend in all three spatial dimensions – a remarkable instrument!

Now consider multiple GW sources. At the low GW frequencies probed by PTAs (where the expected GW point sources are supermassive black hole binaries, far from final merger) the frequency of each gravitational plane wave drifts negligibly over the observation timescale $T \sim 10$ yrs; and over the light travel time from the pulsars to the Earth, the frequency drift or "chirp" is approximately linear: $h_{ij}(t,\vec{x}) = \mathrm{Re}\{\hat{h}_{ij}\mathrm{e}^{-2\pi i\chi(\tau)}\}$, where $\chi(\tau) \equiv f_0\tau + \frac{1}{2}\dot{f}\tau^2$, $\tau \equiv t - \hat{n}\cdot\vec{x}$, and $\{\hat{h}_{ij}, f_0, \dot{f}\}$ are constants. The induced timing residuals for pulsar α are a sum of two peaks in frequency space: an "Earth term" at frequency f_0, and a "pulsar term" at frequency $f_0 - \dot{f}r_\alpha(1+\hat{n}\cdot\hat{r}_\alpha)$; if $\dot{f}$ is large enough (*i.e.* for supermassive black hole binaries of sufficiently high mass, sufficiently close to merger) these two peaks may lie in separate frequency bins [6]. The number of such GW sources that may be individually characterized by a PTA may be determined via the following counting argument. To fully specify the pattern of timing residuals, we must provide the following information: in every GW frequency bin, and for each "Earth term" in that bin, we give the associated propagation direction $\hat{n}$, frequency derivative $\dot{f}$, and two complex amplitudes (*i.e.* an the amplitude and phase for both polarization modes), for a total of 7 real numbers. On the other hand, since the angular dependence of Eq. (4) contains spherical harmonics of arbitrarily high angular momentum order, the number of independent measurements collected by the PTA is simply $2N$ per GW frequency bin – namely, the measured amplitude and phase of the timing residuals, for each pulsar, in each frequency bin [16]. To completely characterize the individual sources, the independent measurements must outnumber the parameters to be determined; that is, the PTA can characterize up to an average of $2N/7$ chirping GW point sources per GW frequency bin. For simplicity, this argument neglects "boundary effects" coming from GW sources for which the "Earth term" lies within the detectable frequency range, while the "pulsar term" does not, or vice versa. If we assume that all of the GW sources are monochromatic ($\dot{f} = 0$), the maximum number that can be characterized improves only slightly to $2N/6$ per frequency bin, but the fitting procedure becomes much easier since we can treat each GW frequency bin independently. If the PTA can disentangle and characterize the individual sources, one expects the angular resolution to be diffraction limited $\delta\theta \gtrsim (1/\mathrm{SNR})(\lambda_{gw}/r_{pulsar})$ [the more precise expectation is given by Eq. (14)].

Pulsars with poorly known distances. If the pulsar distance r_α is poorly known relative to $\lambda_{gw}/(1+\hat{n}\cdot\hat{r}_\alpha)$, then $\mathcal{P}_\alpha$ becomes a random phase containing essentially no information; the $\mathcal{A}$ term is washed out, and only the $\mathcal{B}$ term remains in Eq. (12). Such pulsars no longer contribute diffraction-limited information, but all is not lost!

Let us start by giving the key idea, roughly. In spherical coordinates, with the GW source at the north pole and pulsar α at $(\theta_\alpha, \varphi_\alpha)$, the factor $\tilde{h}_{ij}(f)\hat{r}^i_\alpha\hat{r}^j_\alpha/(1+\hat{n}\cdot\hat{r}_\alpha)$ in Eq. (4) is the familiar [15] pattern $\propto \cos 2\varphi_\alpha(1+\cos\theta_\alpha)$, which is singular (since the $\theta_\alpha \to 0$ limit depends on φ_α). Ultimately this singularity is smoothed out by the $[1-\mathcal{P}_\alpha]$ factor in Eq. (4), but this smoothing only kicks in for very small angular separations $\theta_\alpha \lesssim \sqrt{\lambda_{gw}/(2\pi r_\alpha)}$: for such small separations, $\mathcal{P}_\alpha(f_{gw})$ ceases to be a random phase, and $[1-\mathcal{P}_\alpha(f_{gw})]$ vanishes $\propto \theta_\alpha^2$. In other words, when $\lambda_{gw}/(2\pi r_\alpha) \ll 1$, Eq. (4) is *quasi*-singular at $\theta_\alpha = 0$; it becomes *genuinely* singular in the limit $\lambda_{gw}/(2\pi r_\alpha) \to 0$. The fact that $\delta\tilde{t}_\alpha(f)$ varies rapidly (with $\cos 2\varphi_\alpha$ dependence) around a tiny circle of radius $\sqrt{\lambda_{gw}/(2\pi r_\alpha)} \lesssim \theta_\alpha \ll 1$ surrounding the quasi-singularity is the key to localizing the GW source.

To see this in more detail, split the pulsar directions $\hat{r}_\alpha$ into components parallel and perpendicular to the GW source direction $\hat{z}$: $\hat{r}_\alpha = \hat{\rho}_\alpha \sin\theta_\alpha + \hat{z}\cos\theta_\alpha$. Now approach the quasi-singularity in two steps: first consider the "weaker" limit $\sqrt{\lambda_{gw}/(2\pi r_\alpha)} \lesssim \theta_\alpha \ll 1$ in which θ_α is small but $[1-\mathcal{P}_\alpha]$ is *not*; then proceed to the "stronger" limit $\theta_\alpha \ll \sqrt{\lambda_{gw}/(2\pi r_\alpha)} \ll 1$ in which θ_α and $[1-\mathcal{P}_\alpha]$ are *both* small. In the weaker limit, Eq. (4) becomes

$$\delta\tilde{t}_\alpha(f) \approx \frac{i}{\pi f}\tilde{h}_{ij}(f)\hat{\rho}^i_\alpha\hat{\rho}^j_\alpha[1-\mathcal{P}_\alpha(f)] \tag{15}$$

while Eq. (12) becomes

$$\frac{\partial[\delta\tilde{t}_\alpha(f)]}{\delta\gamma^{\bar{\mu}}} = \frac{2i}{\pi f}\frac{C_{\alpha\bar{\mu}}(f)}{\theta_\alpha}[1-\mathcal{P}_\alpha(f)] \tag{16}$$

where

$$C_{\alpha\bar{\mu}}(f) \equiv \tilde{h}_{ij}(f)\hat{\rho}^i_\alpha\big[\epsilon^j{}_{kl}\hat{n}^k\hat{m}^l_{\bar{\mu}} + \hat{\rho}^j_\alpha\hat{\rho}^k_\alpha\hat{n}^l\hat{m}^m_{\bar{\mu}}\epsilon_{klm}\big]. \tag{17}$$

Thus (still in the weaker limit) we have:

$$\frac{\Gamma^\alpha_{\bar{\mu}\bar{\nu}}}{\mathrm{SNR}^2_\alpha} \approx \frac{4}{\theta^2_\alpha}\frac{C_{\alpha\bar{\mu}}(f_{gw})C_{\alpha\bar{\nu}}(f_{gw})^*}{\big|\tilde{h}_{ij}(f_{gw})\hat{\rho}^i_\alpha\hat{\rho}^j_\alpha\big|^2}. \tag{18}$$

Since the second fraction on the right-hand side of this expression is generically $\mathcal{O}(1)$, this says that when a pulsar is near (but not *too* near) a GW source on the sky, its contribution to $\Gamma_{\bar{\mu}\bar{\nu}}$ is typically $\Gamma^\alpha_{\bar{\mu}\bar{\nu}} \sim (4/\theta^2_\alpha)\mathrm{SNR}^2_\alpha$.

In the stronger limit $\theta_\alpha \ll \sqrt{\lambda_{gw}/(2\pi r_\alpha)} \ll 1$, Eqs. (4) and (12) imply that $\delta\tilde{t}_\alpha(f)$ and $\partial[\delta\tilde{t}_\alpha(f)]/\partial\gamma^{\bar{\mu}}$ are smooth and vanishing at $\theta_\alpha = 0$. So as θ_α decreases, $\Gamma^\alpha_{\bar{\mu}\bar{\nu}}$ initially increases as $1/\theta^2_\alpha$, and then drops to zero; in between it attains a maximum value:

$$\Gamma^\alpha_{\bar{\mu}\bar{\nu}} \sim \frac{8\pi r_\alpha}{\lambda_{gw}}\mathrm{SNR}^2_\alpha \tag{19}$$

at a separation angle $\theta_\alpha \approx \sqrt{\lambda_{gw}/(2\pi r_\alpha)} \ll 1$.

Now consider multiple GW sources. We can repeat the previous section's counting argument, except that we must now include the N unknown pulsar distances when we are counting the parameters needed to specify the pattern of timing residuals; with this modification, we find that a PTA which monitors F different GW frequency bins can completely characterize up to an average of $(2N/7)(1-1/2F)$ "chirping" sources [or $(2N/6)(1-1/2F)$ monochromatic sources] per bin. Note

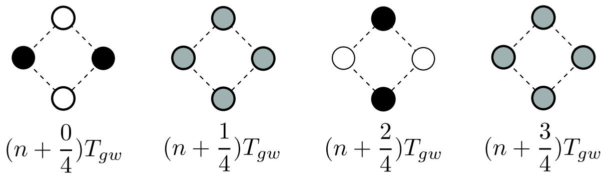

FIG. 1: The 4 circles represent 4 pulsars that form a small square on the sky. The timing residuals of all 4 pulsars are oscillating with the same amplitude and period T_{gw}, but different phases; each pulsar's oscillation is 180^o out of phase with its 2 nearest neighbors. (This figure depicts this by showing 4 different moments in the oscillation cycle: black, white, or grey circles indicate that, at that moment, the pulses are arriving early, late, or "on time," respectively.) This signature indicates that the square contains a GW point source.

that, although we didn't know the pulsar distances *a priori*, they are determined, in principle, by the fit [17] [18].

If the PTA can disentangle and characterize the individual sources, how well can they be angularly localized? To answer this question, one should ask, for each combination of pulsar and GW source, whether the fit to the timing residuals has determined r_α accurately or poorly relative to $\lambda_{gw}/(1+\hat{n}\cdot\hat{r}_\alpha)$; roughly speaking, if r_α has been determined accurately then we expect the pulsar will contribute diffraction limited angular information as described by Eq. (14); and if r_α has been determined poorly then we expect the pulsar will contribute "quasi-singularity limited" angular information $\Gamma^\alpha_{\bar{\mu}\bar{\nu}}$ for that source, as described by Eqs. (18) and (19). Consider the localization of a GW source when all of the pulsar's have poorly known distances; as explained above, the quasi-singular pattern of timing residuals implies that the angular localization will be dominated by the pulsars that are close to that source on the celestial sphere; in particular, it is roughly set by the smallest quadrilateral of pulsars that encircles the source on the celestial sphere, down to a limiting angular resolution of roughly $\delta\theta \sim (1/\mathrm{SNR}_\alpha)\sqrt{\lambda_{gw}/d_{pulsar}}$ [the more precise statement is given by Eqs. (18) and (19)]. To understand this behavior, consider the example in Fig. 1.

Discussion. In the previous sections, we have attempted to clarify the limits on the capabilities of PTAs and, in particular, how these limits depend on factors such as the SNR distribution of the GW sources, the number and angular distribution of the pulsars relative to the GW sources, the distances to the pulsars and the precisions of those distances. Our bounds on the angular resolution were obtained by Fisher matrix methods; as such, these bounds are saturated for GW sources with high SNR, and still provide useful guidelines for modest SNR sources, but will become "loose" for low SNR sources. (Note that the relevant SNR here is the *total* SNR of the source in the PTA, which can be high even if the SNR per pulsar is not.) Our limits on the number of sources that can be individually characterized relied on deterministic counting arguments; again, these limits will be saturated for high SNR sources, and loose for low SNR sources; but even at low SNR, the essential point remains that a PTA with many pulsars can distinguish many sources per frequency bin, so the traditional rule of thumb for confusion (that a gravitational wave detector becomes confused when there is $\gtrsim$ one source per bin) is too pessimistic for PTAs. Interesting problems for future work include: (i) "tightening" the angular resolution and confusion limits at low SNR; (ii) extending this work to GW point sources that are near enough that their wavefront curvature is significant [11]; (iii) determining the circumstances in which pulsar distance determination by GW fitting can compete with more traditional methods (*i.e.* VLBI or timing parallax); (iv) clarifying the statistics of GW sources which are anomalously well characterized because they are fortuitously located relative to one or several pulsars on the sky; (v) quantifying the gain from matched filtering (with quasi-singular filters in particular) compared to traditional stochastic correlation analysis, even when the PTA *appears* source confused.

Acknowledgements. We are grateful to Chris Hirata and Neil Turok for valuable conversations. LB acknowledges support from the CIFAR JFA.

[1] M.V. Sazhin, Soviet Astronomy **22**, 36 (1978).
[2] S. Detweiler, ApJ **234**, 1100 (1979).
[3] R.S. Foster and D.C. Backer, ApJ **361**, 300 (1990).
[4] F. A. Jenet *et al.*, ApJ **653**, 1571 (2006).
[5] R. van Haasteren *et al.*, MNRAS **395**, 1005 (2009).
[6] F. A. Jenet *et al.*, ApJ **606**, 799 (2004).
[7] A. Sesana *et al.*, MNRAS **394**, 2255 (2009).
[8] A. Sesana and A. Vecchio, arXiv:1003.0677.
[9] L. Finn and A. Lommen, ApJ **718**, 1400 (2010).
[10] D. R. Yardley *et al.*, MNRAS **407**, 669 (2010).
[11] X. Deng and L.S. Finn, arXiv:1008:0320.
[12] V. Corbin and N. Cornish, arXiv:1008.1782
[13] K. J. Lee *et al.*, to appear (2010).
[14] C. Misner, K. Thorne and J. Wheeler, *Gravitation*.
[15] R. Hellings and G. Downs, ApJ **361**, 300 (1983).
[16] If Eq. (4) only contained spherical harmonics up to angular momentum order $l \leq l_{max}$, the number of independent measurements collected by the PTA would be limited by $(l_{max}+1)^2$, the total number of spherical harmonics up to this order.
[17] In the special case that the gravitational wave signal is due to a single monochromatic gravitational plane wave, this fit is highly degenerate, since each r_α is only determined modulo $\lambda_{gw}/(1+\hat{n}\cdot\hat{r}_\alpha)$; but, if the GWs are significantly non-monochromatic (*e.g.* "chirping"), significantly non-planar [11] or, most importantly, if there are two or more GW signals propagating in different directions, this degeneracy is broken.
[18] After completion of this work, we learned that Ref. [12] recently emphasized a similar point.