Gravitational waves from inspiraling compact binaries: Validity of the stationary-phase approximation to the Fourier transform

Serge Droz,* Daniel J. Knapp,† and Eric Poisson
Department of Physics, University of Guelph, Guelph, Ontario, Canada N1G 2W1

Benjamin J. Owen
Theoretical Astrophysics 130-33, California Institute of Technology, Pasadena, California 91125
and Max Planck Institut für Gravitationsphysik, Am Mühlenberg 5, 14476 Golm, Germany
(Received 26 January 1999; published 18 May 1999)

We prove that the oft-used stationary-phase method gives a very accurate expression for the Fourier transform of the gravitational-wave signal produced by an inspiraling compact binary. We give three arguments. First, we analytically calculate the next-order correction to the stationary-phase approximation, and show that it is small. This calculation is essentially an application of the steepest-descent method to evaluate integrals. Second, we numerically compare the stationary-phase expression to the results obtained by fast Fourier transform. We show that the differences can be fully attributed to the windowing of the time series, and that they have nothing to do with an intrinsic failure of the stationary-phase method. And third, we show that these differences are negligible for the practical application of matched filtering. [S0556-2821(99)00414-2]

PACS number(s): 04.25.Nx, 04.30.Db, 04.80.Nn

I. INTRODUCTION AND SUMMARY

The ongoing construction of kilometer-size interferometric gravitational-wave detectors, such as the American LIGO (Laser Interferometer Gravitational-wave Observatory), the French-Italian VIRGO, the British-German GEO600, and the Japanese TAMA300, has motivated a lot of recent work on strategies to analyze the data to search for and measure the gravitational-wave signals [1–16]. Much of this work has been devoted to inspiraling compact binaries, composed of neutron stars and/or black holes, which are one of the most promising sources of gravitational waves for these detectors. The key idea behind this work is that the gravitational-wave signals will be detected and measured by matched filtering [17], a well-known technique by which the detector's data stream is cross-correlated with a set of model waveforms (called templates), and the signal-to-noise ratio maximized as a function of the template parameters.

While the detector output is represented by a discrete time series, the operations associated with matched filtering are usually carried out in (discrete) frequency space, which requires taking the (discrete) Fourier transform of the time series; the method of choice here is the standard fast Fourier transform (FFT). At the same time, it is necessary to compute the Fourier transform of $h(t)$, the model waveform. The waveform $h(t)$ can easily be discretized to mimic the discrete sampling of the detector output, and the discrete time series can easily be Fourier transformed by FFT to yield the discrete analogue of the (continuous) Fourier transform $\tilde{h}(f)$.

However, in theoretical investigations [1–13,15,16] it is often much more convenient to deal with $\tilde{h}(f)$ as an analytic expression rather than as a table of values; efforts to obtain such an analytical expression, even if it is only approximate, are therefore well justified. It can also be argued that to express the waveform in the frequency domain is in a sense more natural than to express it in the time domain. One of the reasons is that the orbital energy, whose expression is required in the derivation of the waveform, is primarily a function of F, the gravitational wave's instantaneous frequency (which will defined precisely below); it is therefore natural to express h also as a function of F, which unlike t is a coordinate-independent quantity. Another reason resides with the fact that the post-Newtonian expression for the relation $F(t)$ suggests that the instantaneous frequency is not always a monotonically increasing function of time [18], a nonphysical behavior that signals the eventual breakdown of the post-Newtonian expansion; in contrast, the post-Newtonian expression for the inverted relation $t(F)$ *is* monotonic, a property that suggests that F is indeed a more natural time variable.

In the past, and in the specific context of inspiraling binaries, analytic expressions for $\tilde{h}(f)$ have been obtained within the stationary-phase approximation [19], essentially the leading-order term in an expansion in powers of the small quantity (radiation-reaction time scale)/(orbital period). While it has generally been felt that this approximation is adequate, this belief has not yet been backed up (in the published literature) with a detailed quantitative analysis. Claims in the literature [14] to the effect that this approximation is *not* adequate have prompted us to examine this question.

Our objective in this paper is to prove, beyond any reasonable doubt, that the stationary-phase approximation to the Fourier transform of an inspiraling-binary signal is in fact very accurate. First, we calculate the next-order correction to the stationary-phase approximation, and show that it is indeed small. This calculation is essentially an application of

*Present address: Institute for Theoretical Physics, University of Zurich, CH-8057 Zurich, Switzerland.

†Present address: Department of Physics and Astronomy, McMaster University, Hamilton, Ontario, Canada L8S 4M1.

the steepest-descent method to evaluate integrals [19]. We then numerically compare the stationary-phase expression for $\tilde{h}(f)$ to the results obtained by FFT. We show that the differences we observe can be fully attributed to the windowing of the time series, and that they are irrelevant for matched filtering. Greater differences observed in other cases [14,18] must be attributed not to any intrinsic deficiency of the stationary phase approximation, but rather to one of the other issues in constructing the most accurate expression for $\tilde{h}(f)$—such as the proper (Taylor or Padé) expansion of dF/dt.

We now present and explain our main results. As we shall justify below, it is sufficient for our purposes to consider the leading-order expression for the gravitational-wave signal, obtained by assuming that the binary's orbital motion (taken to be circular) is governed by Newtonian gravity, with an inspiral—caused by the loss of energy and angular momentum to gravitational radiation—governed by the Einstein quadrupole formula. This so-called *Newtonian* signal is given by [1]

$$h(t) = Q(\text{angles}) \frac{\mathcal{M}}{r} (\pi \mathcal{M} F)^{2/3} e^{-i\Phi}, \tag{1.1}$$

where Q is a function of all the relevant angles (position of the source in the sky, inclination of the orbital plane, orientation of the gravitational-wave detector), $\mathcal{M} = (m_1 m_2)^{3/5}/(m_1+m_2)^{1/5}$ (with m_1 and m_2 denoting the individual masses) is the *chirp mass*, r is the distance to the source, $F(t)$ the instantaneous frequency (twice the orbital frequency), and $\Phi(t) = \int 2\pi F(t')\, dt'$ is the phase. The function $h(t)$ represents the gravitational wave measured at the detector site. It is a linear superposition of the wave's two fundamental polarizations, and we choose to express it in this complex form for convenience. The relation between F and t is given by [1]

$$\pi \mathcal{M} F(t) = \left[\frac{5\mathcal{M}}{256(t_c - t)} \right]^{3/8}, \tag{1.2}$$

where t_c, the ''time at coalescence,'' is such that formally, $F(t_c) = \infty$. The relation between Φ and t is

$$\Phi(t) = \Phi_c - \frac{1}{16} \left[\frac{256(t_c - t)}{5\mathcal{M}} \right]^{5/8}, \tag{1.3}$$

where Φ_c, the ''phase at coalescence,'' is equal to $\Phi(t_c)$. The Newtonian signal is the leading-order term in the expansion of the gravitational waves in powers of $V \ll 1$, where $V = (\pi \mathcal{M} F)^{1/3}$ is (up to a numerical factor) the orbital velocity. Post-Newtonian corrections to this result come with a relative factor of order V in the amplitude, and a relative factor of order V^2 in the phase [20]. Throughout this paper we use geometrized units, setting $G = c = 1$.

In the stationary-phase approximation, the Fourier transform of the function $h(t)$ appearing in Eq. (1.1) is given by [5]

$$\tilde{h}_{\text{spa}}(f) = \frac{\sqrt{30\pi}}{24} \frac{Q\mathcal{M}^2}{r} v^{-7/2} e^{i\psi}, \tag{1.4}$$

where $v \equiv (\pi \mathcal{M} f)^{1/3} \ll 1$, and

$$\psi(v) = 2\pi f t_c - \Phi_c - \frac{\pi}{4} + \frac{3}{128 v^5}. \tag{1.5}$$

Two types of corrections to this result are calculated in Sec. II.

The first type of correction constitutes an intrinsic improvement on the stationary-phase approximation. We show that the steepest-descent evaluation of the Fourier transform returns the same amplitude as before, but that the phase is altered by a term $\delta\psi$: $\psi \to \psi + \delta\psi$, where

$$\delta\psi(v) = \frac{92}{45} v^5 + O(v^{10}). \tag{1.6}$$

Notice that this is a correction of order v^{10} *relative* to $\psi(v)$. This is much smaller than post-Newtonian corrections to the phase, which appear at relative order v^2 [5]. Incorporating these post-Newtonian corrections into our calculation would only change Eq. (1.6) by adding a term of order v^7 to the right-hand side. This justifies the fact that it was not necessary, for the purposes of this investigation, to use more accurate versions of Eqs. (1.2) and (1.3).

The second type of correction addresses an implicit assumption of the stationary-phase method, that the function $h(t)$ has support in the complete time interval $-\infty < t < t_c$. Physically, this assumption means that the binary system must have formed in the infinite past (a reasonable assumption given the long lives of compact binaries), and that the inspiral must continue until $F = F(t_c) = \infty$ (an unrealistic assumption). If we choose instead to restrict the time interval to $t_{\min} < t < t_{\max}$, such that $F_{\min} \equiv F(t_{\min}) > 0$ and $F_{\max} \equiv F(t_{\max}) < \infty$, then the Fourier transform will be affected, and it will differ from $\tilde{h}_{\text{spa}}(f)$. The value of $F_{\min}$ is typically chosen to reflect the lower bound of the instrument's frequency band. The value of $F_{\max}$ could be chosen to reflect the upper bound of the instrument's frequency band, or for some binaries it can be chosen to correspond to the approximate frequency of the last stable orbit, at which the inspiral signal changes over to a poorly-known merger signal.

We shall refer to this truncation of the time interval as ''windowing,'' and for concreteness, we will assume that the signal is started abruptly at $t = t_{\min}$ and ended abruptly at $t = t_{\max}$. Thus, $h(t)$ is assumed to be given by Eq. (1.1) in the interval $t_{\min} < t < t_{\max}$, and is assumed to be zero outside this interval. In Sec. II we show that windowing affects both the amplitude and the phase of the Fourier transform. The amplitude acquires an extra factor $1 + \delta A_{\text{w}}$, where

$$\delta A_{\text{w}}(v) = -\frac{12}{\sqrt{30\pi}} v^{7/2} \left[\frac{{x_{\min}}^2}{v^3 - {x_{\min}}^3} \cos\left(\phi_{\min} + \frac{\pi}{4} \right) + \frac{{x_{\max}}^2}{{x_{\max}}^3 - v^3} \cos\left(\phi_{\max} + \frac{\pi}{4} \right) \right], \tag{1.7}$$

where $x_{\min}\equiv(\pi\mathcal{M}F_{\min})^{1/3}$ and

$$\phi_{\min}=\frac{5v^3-8{x_{\min}}^3}{128{x_{\min}}^8}+\frac{3}{128v^5}, \tag{1.8}$$

with similar equations holding for $x_{\max}$ and $\phi_{\max}$. On the other hand, the phase acquires an extra term $\delta\psi_{\rm w}$ given by

$$\delta\psi_{\rm w}(v)=\frac{12}{\sqrt{30\pi}}v^{7/2}\left[\frac{{x_{\min}}^2}{v^3-{x_{\min}}^3}\sin\left(\phi_{\min}+\frac{\pi}{4}\right)+\frac{{x_{\max}}^2}{{x_{\max}}^3-v^3}\sin\left(\phi_{\max}+\frac{\pi}{4}\right)\right]. \tag{1.9}$$

The fact that $\delta A_{\rm w}$ and $\delta\psi_{\rm w}$ both diverge at the boundaries $v=x_{\min}$ and $v=x_{\max}$ signals the breakdown of our approximations there. Away from the boundaries, $\delta A_{\rm w}$ and $\delta\psi_{\rm w}$ are bounded, and they oscillate as a function of frequency. Thus, these corrections represent amplitude and phase modulations induced by the abrupt cutoffs of the function $h(t)$ at the boundary points. This is to be contrasted with $\delta\psi$, which represents a steady phase drift.

Our complete expression for the steepest-descent approximation to the Fourier transform is therefore

$$\tilde{h}_{\rm sda}(f)=\tilde{h}_{\rm spa}(f)(1+\delta A_{\rm w})e^{i(\delta\psi+\delta\psi_{\rm w})}. \tag{1.10}$$

In Sec. III we show that the only noticeable corrections to $\tilde{h}_{\rm spa}(f)$ are the amplitude and phase modulations that come as a consequence of windowing; in particular, the intrinsic correction $\delta\psi$ is too small to be noticeable in the relevant frequency interval. We do this by comparing $\tilde{h}_{\rm spa}(f)$ to $\tilde{h}_{\rm fft}(f)$, the discrete Fourier transform of the windowed time series $h(t)$; this comparison reveals that any discrepancy between the two versions of the Fourier transform can be fully accounted for by the modulations $\delta A_{\rm w}$ and $\delta\psi_{\rm w}$. This allows us to conclude that windowing, and windowing only, must be held responsible for any discrepancy between $\tilde{h}_{\rm spa}(f)$ and $\tilde{h}_{\rm fft}(f)$. While limited to the Newtonian signal of Eq. (1.1), there is no reason to believe that this conclusion would be invalidated by a full post-Newtonian analysis [21].

Finally, in Sec. IV we calculate the matched-filtering overlap between $\tilde{h}_{\rm spa}(f)$ and $\tilde{h}_{\rm fft}(f)$, and show that the modulations do not significantly affect the overlap. This result, together with our previous findings, lead us to conclude that for the purposes of matched filtering, $\tilde{h}_{\rm spa}(f)$ and $\tilde{h}_{\rm fft}(f)$ are essentially *equivalent* representations of the gravitational-wave signal.

The main conclusion of this work is that the stationary-phase method returns a sufficiently accurate expression for $\tilde{h}(f)$; a similar conclusion was reached independently by Chassande-Mottin and Flandrin, and reported in Refs. [22,23]. What is more, from the fact that $\delta\psi/\psi=O(v^{10})$, we can be sure that the method will stay accurate for as long as the post-Newtonian expansion of $h(t)$ in powers of V is itself an accurate approximation to the gravitational-wave signal.

II. CALCULATION OF THE FOURIER TRANSFORM

We begin with the time-domain signal of Eq. (1.1), and we assume that the signal begins abruptly at a time $t_{\min}$ and ends abruptly at a time $t_{\max}$. We let $F_{\min}\equiv F(t_{\min})$ and $F_{\max}\equiv F(t_{\max})$ be the corresponding instantaneous frequencies. The relation between F and t is obtained by integrating

$$\frac{dF}{dt}=\frac{96}{5\pi\mathcal{M}^2}(\pi\mathcal{M}F)^{11/3}, \tag{2.1}$$

which leads to

$$t(F)=t_c-\frac{5\mathcal{M}}{256}(\pi\mathcal{M}F)^{-8/3}, \tag{2.2}$$

where t_c (''time at coalescence'') is a constant of integration. The phase function is then given by

$$\Phi(F)=\Phi_c-\frac{1}{16}(\pi\mathcal{M}F)^{-5/3}, \tag{2.3}$$

where Φ_c (''phase at coalescence'') is another constant of integration.

The Fourier transform,

$$\tilde{h}(f)=\int h(t)e^{2\pi ift}\,dt, \tag{2.4}$$

is evaluated by introducing a new integration variable,

$$x\equiv(\pi\mathcal{M}F)^{1/3}, \tag{2.5}$$

which can be related to t via Eqs. (2.1) and (2.2). After some rearrangement, we obtain

$$\tilde{h}(f)=\frac{5Q\mathcal{M}^2}{32r}\exp\left[i\left(2\pi ft_c-\Phi_c+\frac{3}{128v^5}\right)\right]\times I(v,x_{\min},x_{\max}), \tag{2.6}$$

where $x_{\min}=(\pi\mathcal{M}F_{\min})^{1/3}$, $x_{\max}=(\pi\mathcal{M}F_{\max})^{1/3}$, and

$$v\equiv(\pi\mathcal{M}f)^{1/3}\ll 1. \tag{2.7}$$

We have introduced the Fourier integral

$$I(v,x_{\min},x_{\max})=\int_{x_{\min}}^{x_{\max}}x^{-7}e^{-i\phi}\,dx, \tag{2.8}$$

where

$$\phi(x;v)=\frac{5v^3-8x^3}{128x^8}+\frac{3}{128v^5}. \tag{2.9}$$

It is easy to check that $\phi(v;v)=\phi'(v;v)=0$, where a prime denotes differentiation with respect to x. The function $\phi(x;v)$ initially decreases from $\phi_{\min}(v)\equiv\phi(x_{\min};v)$ to zero as x increases from $x_{\min}$ to v, and then increases from zero to $\phi_{\max}(v)\equiv\phi(x_{\max};v)$ as x increases from v to $x_{\max}$.

The Fourier integral is evaluated by using ϕ as the integration variable. Because ϕ is not monotonic in the interval $[x_{\min},x_{\max}]$, the integral must be broken down into two parts. The first part covers the interval $[x_{\min},v)$, while the second part covers the interval $(v,x_{\max}]$. It is easy to check that the integral can be expressed as

$$I(v,x_{\min},x_{\max})=\frac{16}{5}[J_1(v,x_{\min})+J_2(v,x_{\max})], \tag{2.10}$$

where

$$J_1(v,x_{\min})=\int_0^{\phi_{\min}}\frac{x^2}{v^3-x^3}e^{-i\phi}\,d\phi \tag{2.11}$$

and

$$J_2(v,x_{\max})=\int_0^{\phi_{\max}}\frac{x^2}{x^3-v^3}e^{-i\phi}\,d\phi. \tag{2.12}$$

We recall that $\phi_{\min}=\phi(x_{\min};v)$ and $\phi_{\max}=\phi(x_{\max};v)$ are functions of v.

The coordinate transformation $x\to\phi$ is defined implicitly by Eq. (2.9). It will prove sufficient to invert this relation in a neighborhood of $x=v$, or $\phi=0$. The following relations are established by Taylor expansion:

$$x=v\mp\frac{4\sqrt{30}}{15}v^{7/2}\phi^{1/2}+\frac{256}{45}v^6\phi\mp\frac{10016\sqrt{30}}{2025}v^{17/2}\phi^{3/2}+\frac{4281344}{30375}v^{11}\phi^2+O(v^{27/2}\phi^{3/2}), \tag{2.13}$$

where the upper sign refers to the interval $[x_{\min},v)$, while the lower sign refers to $(v,x_{\max}]$. We also have

$$\frac{x^2}{v^3-x^3}=\frac{\sqrt{30}}{24}v^{-7/2}\phi^{-1/2}+\frac{5}{9}v^{-1}-\frac{23\sqrt{30}}{135}v^{3/2}\phi^{1/2}+\frac{19232}{6075}v^4\phi+O(v^{13/2}\phi^{3/2}) \tag{2.14}$$

in the first interval, and

$$\frac{x^2}{x^3-v^3}=\frac{\sqrt{30}}{24}v^{-7/2}\phi^{-1/2}-\frac{5}{9}v^{-1}-\frac{23\sqrt{30}}{135}v^{3/2}\phi^{1/2}-\frac{19232}{6075}v^4\phi+O(v^{13/2}\phi^{3/2}) \tag{2.15}$$

in the second interval.

To evaluate the integrals $J_1(v,x_{\min})$ and $J_2(v,x_{\max})$, we let $\phi=\alpha-i\beta$ and deform the contour of integration into the complex plane. While the original contour is along the α axis, from 0 to $\phi_{\min}$ or $\phi_{\max}$, we take the new contour to be the union of C and C', where C is the curve $\alpha=0$ with β running from 0 to ∞, while C' is the curve $\alpha=\phi_{\min}$ or $\phi_{\max}$ with β running from ∞ back to 0. The contour is completed by joining C and C' with the curve $\beta=\infty$, with α running from 0 to $\phi_{\min}$ or $\phi_{\max}$; this part of the contour does not contribute to the integral. The advantage of this choice of contour is that the integrand is exponentially suppressed away from $\beta=0$, ensuring a rapid convergence of the integral.

We evaluate the contribution from C to $J_1(v,x_{\min})$ by substituting Eq. (2.14) into Eq. (2.11), replacing ϕ by $-i\beta$ and using $\beta=0$ and $\beta=\infty$ as limits. The integrations give rise to Γ-functions, and we obtain

$$J_1^C=\frac{\sqrt{-30i\pi}}{24}v^{-7/2}-\frac{5i}{9}v^{-1}+\frac{23\sqrt{30i\pi}}{270}v^{3/2}-\frac{19232}{6075}v^4+O(v^{13/2}). \tag{2.16}$$

A similar calculation also reveals

$$J_2^C=\frac{\sqrt{-30i\pi}}{24}v^{-7/2}+\frac{5i}{9}v^{-1}+\frac{23\sqrt{30i\pi}}{270}v^{3/2}+\frac{19232}{6075}v^4+O(v^{13/2}). \tag{2.17}$$

The contribution from C' to $J_1(v,x_{\min})$ is calculated by letting $\phi=\phi_{\min}-i\beta$, and expressing the function $f(\phi)\equiv x^2/(v^3-x^3)$ as a Taylor series about $\phi_{\min}$. Thus, $f(\phi_{\min}-i\beta)=f(\phi_{\min})-if'(\phi_{\min})\beta+\cdots$ is substituted into Eq. (2.11), whose limits are replaced by $\beta=\infty$ and $\beta=0$. The resulting integrations are again elementary, and we obtain

$$J_1^{C'}=ie^{-i\phi_{\min}}g(v,x_{\min}), \tag{2.18}$$

where

$$g(v,x)=\frac{x^2}{|v^3-x^3|}+\frac{16i}{5}\,\frac{x^{10}(x^3+2v^3)}{|v^3-x^3|^3}+\cdots. \tag{2.19}$$

A similar calculation also reveals

$$J_2^{C'}=ie^{-i\phi_{\max}}g(v,x_{\max}). \tag{2.20}$$

We note that when $x\ll 1$, the function $g(v,x)$ can be well approximated by its first term, $x^2/|v^3-x^3|$, except when $x\simeq v$. In this situation, the expansion for $g(v,x)$ does not converge, and our method of calculation breaks down. Thus, our expression for the Fourier transform will be accurate only when $v^3\equiv\pi\mathcal{M}f$ is not too close to either ${x_{\min}}^3\equiv\pi\mathcal{M}F_{\min}$ or ${x_{\max}}^3\equiv\pi\mathcal{M}F_{\max}$. In other words, our expression will be inaccurate near the boundaries $f=F_{\min}$ and $f=F_{\max}$.

It is easy to see from Eqs. (2.18) and (2.19) that the contribution from C' to $J_1(v,x_{\min})$ vanishes in the limit $x_{\min}\to 0$. Similarly, it can be shown that the contribution from C' to $J_2(v,x_{\max})$ vanishes (as ${x_{\max}}^{-6}$) in the limit $x_{\max}\to\infty$. [This behavior is not revealed by Eq. (2.19), which gives $g(v,x)$ as a series expansion for small values of x. An alternative expression for $g(v,x)$, appropriate for large values of

x, can easily be obtained.] This allows us to conclude that if the boundaries are pushed to $F_{\min}=0$ and $F_{\max}=\infty$, then C' no longer contributes to $J_1(v,x_{\min})$ and $J_2(v,x_{\max})$.

Gathering the results, Eqs. (2.10) and (2.16)–(2.20), we find that the Fourier integral can be expressed as

$$I(v,x_{\min},x_{\max})=\frac{4\sqrt{30\pi}}{15}e^{-i\pi/4}v^{-7/2}\times\left[1+\frac{92}{45}iv^5+O(v^{10})+R(v,x_{\min},x_{\max})\right], \tag{2.21}$$

where

$$R(v,x_{\min},x_{\max})=-\frac{12}{\sqrt{30\pi}}e^{-i\pi/4}v^{7/2}[e^{-i\phi_{\min}}g(v,x_{\min})+e^{-i\phi_{\max}}g(v,x_{\max})]. \tag{2.22}$$

Substituting this into Eq. (2.6), treating R and the $O(v^5)$ term in Eq. (2.21) as small quantities, we arrive at the following expression for the Fourier transform:

$$\tilde{h}(f)=\tilde{h}_{\rm spa}(f)(1+\delta A_{\rm w})e^{i(\delta\psi+\delta\psi_{\rm w})}. \tag{2.23}$$

Here,

$$\tilde{h}_{\rm spa}(f)=\frac{\sqrt{30\pi}}{24}\frac{Q\mathcal{M}^2}{r}v^{-7/2}e^{i\psi}, \tag{2.24}$$

with

$$\psi(v)=2\pi ft_c-\Phi_c-\frac{\pi}{4}+\frac{3}{128v^5}, \tag{2.25}$$

is the stationary-phase approximation to the Fourier transform. Our calculation reveals the existence of two types of correction terms. The first is

$$\delta\psi(v)=\frac{92}{45}v^5+O(v^{10}), \tag{2.26}$$

which represents a small, but steadily growing phase drift. Notice that $\delta\psi$ is of order $O(v^{10})$ relative to ψ. The other correction terms come as a consequence of the abrupt cutoffs imposed at $F=F_{\min}$ and $F=F_{\max}$. They are

$$\delta A_{\rm w}(v)=-\frac{12}{\sqrt{30\pi}}v^{7/2}\left[g(v,x_{\min})\cos\left(\phi_{\min}+\frac{\pi}{4}\right)+g(v,x_{\max})\cos\left(\phi_{\max}+\frac{\pi}{4}\right)\right] \tag{2.27}$$

and

$$\delta\psi_{\rm w}(v)=\frac{12}{\sqrt{30\pi}}v^{7/2}\left[g(v,x_{\min})\sin\left(\phi_{\min}+\frac{\pi}{4}\right)+g(v,x_{\max})\sin\left(\phi_{\max}+\frac{\pi}{4}\right)\right]. \tag{2.28}$$

Notice that $\delta A_{\rm w}$ represents an amplitude modulation, while $\delta\phi_{\rm w}$ is a phase modulation; both oscillate as a function of frequency. The suffix ''w'' indicates that these corrections are associated with ''windowing.''

III. COMPARISON WITH DISCRETE FOURIER TRANSFORM

The preceding analysis reveals that apart from windowing issues, the stationary-phase approximation to the Fourier transform is extremely accurate: Apart from the modulations $\delta A_{\rm w}$ and $\delta\psi_{\rm w}$, $\tilde{h}_{\rm spa}(f)$ differs from $\tilde{h}(f)$ only by a small phase drift $\delta\psi$ of relative order v^{10}. In this section, we firm up this conclusion by comparing $\tilde{h}_{\rm spa}(f)$ to $\tilde{h}_{\rm fft}(f)$, the discrete Fourier transform of the function $h(t)$.

The discrete Fourier transform is evaluated by fast Fourier transform (FFT), using the routines of *Numerical Recipes* [24]. The time series is prepared as follows.

We begin with $h(t)$ as displayed in Eq. (1.1), with the irrelevant factor $Q\mathcal{M}/r$ set to unity. Thus,

$$h(t)=(\pi\mathcal{M}F)^{2/3}e^{-i\Phi}, \tag{3.1}$$

where $F(t)$ and $\Phi(t)$ are given by Eqs. (1.2) and (1.3). This function is assumed to be nonzero only in the interval $F_{\min}<F(t)<F_{\max}$. The duration of the signal is

$$T=t(F_{\max})-t(F_{\min}), \tag{3.2}$$

while the total number of wave cycles is

$$\mathcal{N}=\frac{1}{2\pi}[\Phi(F_{\max})-\Phi(F_{\min})]. \tag{3.3}$$

The values of $h(t)$ at the endpoints $t(F_{\min})$ and $t(F_{\max})$ do not agree. This is a potential difficulty for the FFT, which considers the signal to be periodic with period T. To remedy this, we prepare our time series by padding $h(t)$ with zeros on both sides. More precisely, we let $h(t)$ be zero in the interval $0<t<T$, be equal to the expression (3.1) in the interval $T<t<2T$, and be zero again in the interval $2T<t<4T$. Thus, the effective duration of the time series is four times the duration of the actual signal. We choose the value of the parameters t_c and Φ_c such that $t(F_{\min})\equiv T$ and $\Phi(F_{\min})\equiv 0$. This particular padding of the time series is somewhat ad hoc (it has not been carefully chosen to be the smallest needed to make negligible the circular correlations of the FFT), but as we shall see the final results are very good and optimization would be redundant.

The zero-padded function $h(t)$ is discretely sampled at times $t_k=k\Delta t$, where $k=0,1,\ldots,4N-1$ and $\Delta t=4T/(4N-1)$, with $4N$ denoting the number of sampled points. The Nyquist frequency [24] is given by $f_{\rm Ny}=1/(2\Delta t)$, and N

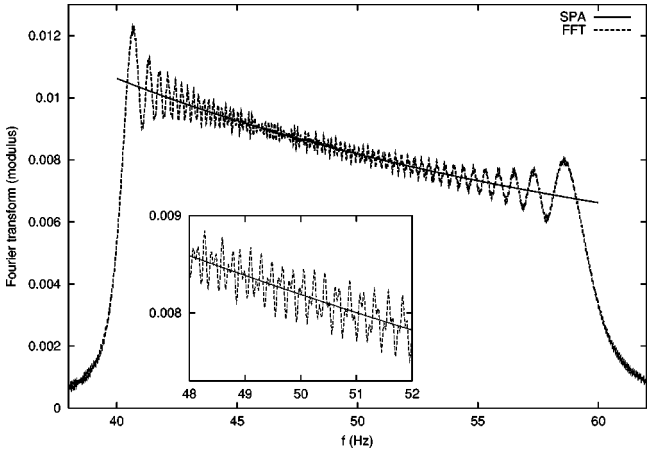

FIG. 1. The solid curve labeled ''SPA'' is a plot of $\mathrm{mod}[\tilde{h}_{\mathrm{spa}}(f)]$, and the dashed curve labeled ''FFT'' is a plot of $\mathrm{mod}[\tilde{h}_{\mathrm{fft}}(f)]$. The inset shows the same curves in a smaller frequency interval.

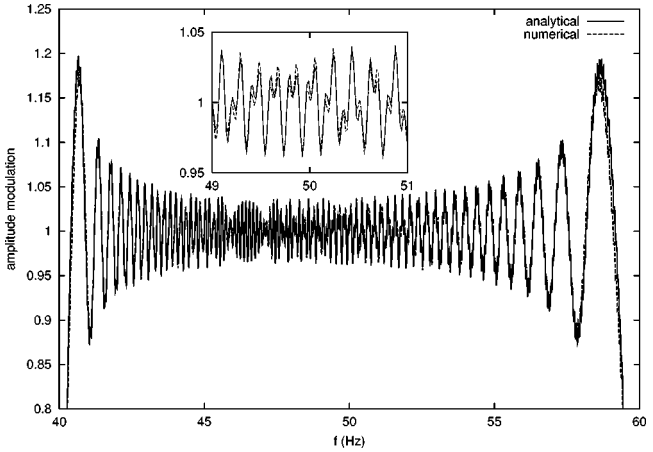

FIG. 2. The solid curve labeled ''analytical'' is a plot of $1+\delta A_{\mathrm{w}}(f)$, and the dashed curve labeled ''numerical'' is a plot of $A_{\mathrm{rel}}(f)$. The inset shows the same curves in a smaller frequency interval.

must be adjusted so that $f_{\mathrm{Ny}}>F_{\mathrm{max}}$. The FFT returns the Fourier transform of the time series, discretely sampled in the frequency domain. The frequency resolution is $\Delta f=(4N\Delta t)^{-1}$.

We denote the discrete Fourier transform of the zero-padded time series $h(t)$ by $\tilde{h}_{\mathrm{fft}}(f)$, and we wish to compare this to $\tilde{h}_{\mathrm{spa}}(f)$, the stationary-phase approximation given by Eqs. (2.24) and (2.25). (It is understood that the correct values for t_c and Φ_c are substituted in these equations.) To do this we define the relative amplitude A_{rel} and relative phase ψ_{rel} by

$$A_{\mathrm{rel}}(f)=\mathrm{mod}[\tilde{h}_{\mathrm{fft}}(f)/\tilde{h}_{\mathrm{spa}}(f)] \tag{3.4}$$

and

$$\psi_{\mathrm{rel}}(f)=\arg[\tilde{h}_{\mathrm{fft}}(f)/\tilde{h}_{\mathrm{spa}}(f)], \tag{3.5}$$

where $\mathrm{mod}(z)=r$ is the modulus of the complex number $z=re^{i\theta}$, while $\arg(z)=\theta$ is its argument. If $\tilde{h}_{\mathrm{fft}}(f)$ and $\tilde{h}_{\mathrm{spa}}(f)$ were in perfect agreement, then $A_{\mathrm{rel}}=1$ and $\psi_{\mathrm{rel}}=0$.

Figure 1 shows plots of $\mathrm{mod}(\tilde{h}_{\mathrm{fft}})$ and $\mathrm{mod}(\tilde{h}_{\mathrm{spa}})$ as functions of f for a signal prepared with $\mathcal{M}=1.25\,M_{\odot}$ in the interval between $F_{\mathrm{min}}=40$ Hz and $F_{\mathrm{max}}=60$ Hz. (Such a narrow band is not at all typical for inspiral signals sought by any interferometer; we use it merely to exaggerate the errors caused by the stationary phase approximation to the level where they are visible.) The duration of such a signal is $T=15.8$ s, and the total number of wave cycles is $\mathcal{N}=749$. The FFT was taken with $N=2048$, giving a Nyquist frequency of $f_{\mathrm{Ny}}=65$ Hz. The figure shows that the agreement is not perfect: While the stationary-phase approximation seems to give the mean curve, the discrete Fourier transform displays oscillations about the mean, and these grow near the boundary points, $f=F_{\mathrm{min}}$ and $f=F_{\mathrm{max}}$. We shall argue that the discrepancy is entirely caused by windowing.

In Fig. 2 we show plots of $A_{\mathrm{rel}}(f)$—defined in Eq. (3.4)—and $1+\delta A_{\mathrm{w}}(f)$—defined in Eq. (1.7). We recall that $\delta A_{\mathrm{w}}(f)$ represents the amplitude modulation induced by windowing. The near-perfect agreement between $A_{\mathrm{rel}}(f)$ and $1+\delta A_{\mathrm{w}}(f)$ shows that any discrepancy between the discrete Fourier transform and the stationary-phase approximation must be attributed to windowing. This conclusion is confirmed by Fig. 3, which shows plots of $\psi_{\mathrm{rel}}(f)$—defined in Eq. (3.5)—and $\delta\psi_{\mathrm{w}}(f)$—defined in Eq. (1.9). Again we see a near-perfect agreement, indicating that windowing accounts fully for any discrepancy between $\tilde{h}_{\mathrm{spa}}(f)$ and $\tilde{h}_{\mathrm{fft}}(f)$.

Equations (1.7) and (1.9) give approximate expressions for $\delta A_{\mathrm{w}}(f)$ and $\delta\psi_{\mathrm{w}}(f)$, and we should expect that in some situations, there could be noticeable differences between these quantities and the numerically-determined $A_{\mathrm{rel}}(f)$ and $\psi_{\mathrm{rel}}(f)$. Figure 4 indicates that such is indeed the case when the frequency interval is expanded. Here, the signal is prepared with the same chirp mass as before, but the frequency interval is now between $F_{\mathrm{min}}=40$ Hz and $F_{\mathrm{max}}=1300$ Hz;

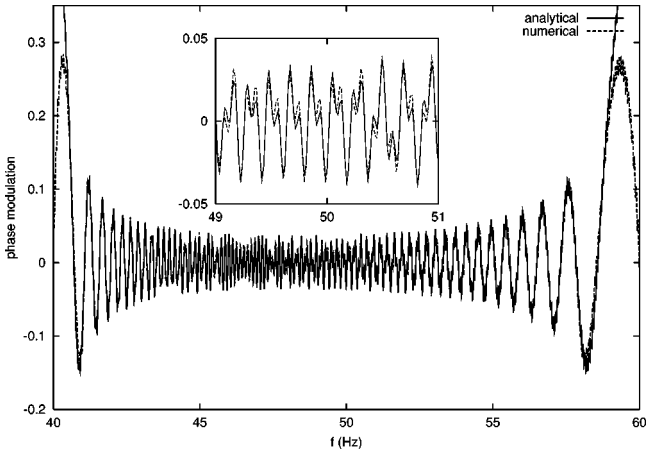

FIG. 3. The solid curve labeled ''analytical'' is a plot of $\delta\psi_{\mathrm{w}}(f)$, and the dashed curve labeled ''numerical'' is a plot of $\psi_{\mathrm{rel}}(f)$. The inset shows the same curves in a smaller frequency interval. Notice that the error in the phase does not accumulate, and that it is always much smaller than 2π.

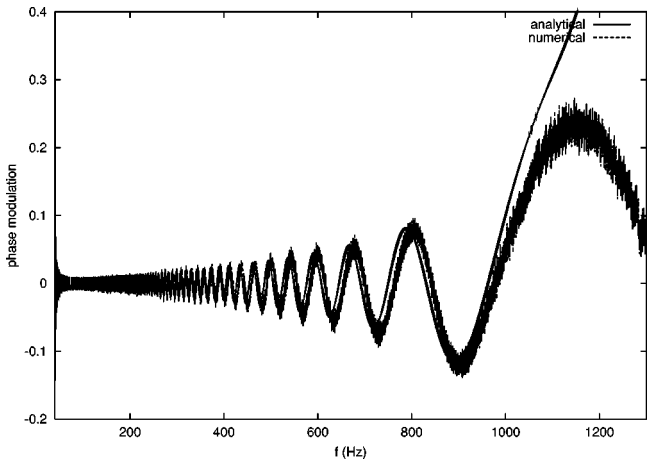

FIG. 4. The thin, solid curve labeled ''analytical'' is a plot of $\delta\psi_{\rm w}(f)$, and the thick, dashed curve labeled ''numerical'' is a plot of $\psi_{\rm rel}(f)$. Notice that here also the error in the phase does not accumulate, and is always much smaller than 2π.

other relevant quantities are listed in Table I below. Although the agreement is no longer near-perfect, it is still remarkably good, and this re-enforces our claim that any discrepancy between the discrete Fourier transform and the stationary-phase approximation is entirely an artifact of windowing. We have verified that the intrinsic correction to the stationary-phase approximation, $\delta\psi(f)$ given by Eq. (1.6), is irrelevant in this frequency interval: This phase shift is just too small to be noticeable at these frequencies.

TABLE I. The last column gives the overlap $\mathcal{O}$ between stationary phase and FFT waveforms prepared with (identical) chirp mass $\mathcal{M}$ given in solar masses in the first column, initial frequency $F_{\rm min}$ given in Hz in the second column, and final frequency $F_{\rm max}$ given in Hz in the third column. The waveforms have a duration T given in seconds in the fourth column and a number of wave cycles $\mathcal{N}$ given in the fifth column. The number of sampled times is $4N$, where N is given in the sixth column, corresponding to a Nyquist frequency $f_{\rm Ny}$ given in Hz in the seventh column. In all cases the value of $F_{\rm max}$ is sufficiently large that the contribution to the overlap from higher frequencies can be neglected (assuming the initial LIGO noise spectrum given in the text).

$\mathcal{M}$ $M_\odot$	$F_{\rm min}$ Hz	$F_{\rm max}$ Hz	T s	$\mathcal{N}$	N	$f_{\rm Ny}$ Hz	$\mathcal{O}$
1.00	40	900	34.6	2200	2^{16}	948	0.9999
1.25	40	1300	23.8	1521	2^{16}	1375	0.9998
1.50	40	1300	17.6	1122	2^{16}	1863	0.9997
1.75	40	1200	13.6	868	2^{15}	1204	0.9997
2.00	40	1500	10.9	695	2^{15}	1505	0.9996
2.25	40	900	8.9	570	2^{14}	916	0.9994
2.50	40	1000	7.5	478	2^{14}	1091	0.9994
2.75	40	1200	6.4	409	2^{14}	1279	0.9994
3.00	40	1400	5.5	354	2^{14}	1479	0.9994
10.00	40	1300	0.7	48	2^{11}	1374	0.9972

IV. OVERLAP INTEGRAL

We have shown that any discrepancy between $\tilde{h}_{\rm fft}(f)$ and $\tilde{h}_{\rm spa}(f)$ can be fully attributed to windowing, which induces amplitude and phase modulations in the Fourier transform. In this section we show that these modulations have no significant effect on operations associated with matched filtering.

The standard theory of matched filtering predicts that the loss of signal-to-noise ratio incurred when filtering a signal $\tilde{h}_{\rm fft}(f)$ with a filter $\tilde{h}_{\rm spa}(f)$ is equal to [8]

$$\mathcal{O}=\frac{(h_{\rm fft}|h_{\rm spa})}{\sqrt{(h_{\rm fft}|h_{\rm fft})(h_{\rm spa}|h_{\rm spa})}}. \tag{4.1}$$

We will refer to this quantity as the *overlap* between the two expressions for the signal's Fourier transform. An overlap close to unity indicates that the filter is an accurate representation of the signal, and that using this filter in analyzing the data will return the largest possible signal-to-noise ratio. We use the notation

$$(a|b)=2\int_{F_{\rm min}}^{F_{\rm max}}\frac{\tilde{a}^*(f)\tilde{b}(f)+\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)}df, \tag{4.2}$$

where $S_n(f)$ is the spectral density of the detector noise.

Notice that the point of view expressed here is that $\tilde{h}_{\rm fft}(f)$ is an exact representation of the signal's Fourier transform, while $\tilde{h}_{\rm spa}(f)$ is an approximate filter. However, because of its symmetry in these quantities, $\mathcal{O}$ also represents the loss in signal-to-noise ratio incurred when filtering a signal $\tilde{h}_{\rm spa}(f)$ with a filter $\tilde{h}_{\rm fft}(f)$. This is the opposite point view, in which the stationary-phase approximation is viewed as an exact representation of the Fourier transform.

We evaluate the integrals in Eq. (4.1) by turning them into discrete sums, using the sampled frequencies $f_k=k\Delta f$ returned by the FFT. Thus, $\int\alpha(f)\,df\approx\Sigma_k\alpha(f_k)\Delta f$. It is sufficient for our purposes to use a simple analytic model for the noise's spectral density. We choose a noise curve that roughly mimics the expected noise spectrum of the initial LIGO detector, and set [5]

$$S_n(f)=S_0[(f_0/f)^4+2+2(f/f_0)^2] \tag{4.3}$$

for $f>40$ Hz, with $f_0=200$ Hz. The value of S_0 is irrelevant for our purposes, and $S_n(f)$ is taken to be infinite below 40 Hz.

The overlap $\mathcal{O}$ is calculated for a number of chirp masses. The results are displayed in Table I. The conclusion is clear: For all cases, $\mathcal{O}>0.997$, indicating that the amplitude and phase modulations have very little effect on matched filtering. In view of the fact that the modulations oscillate and never get large, especially away from $F_{\rm min}$ and $F_{\rm max}$ where the instrument is most sensitive, this is the expected conclusion.

It should be noted that for most of the binaries listed in Table I, the adopted value for $F_{\rm max}$ *exceeds* the frequency at which the last stable orbit is expected to be found. [For equal-mass systems, this frequency is given approximately

by $1910(M_\odot/\mathcal{M})$ Hz.] Our gravitational-wave signals are therefore not realistic at high frequencies, something which is already made clear by the fact that we do not incorporate post-Newtonian corrections into our waveforms. Since our purpose here is simply to establish the validity of the stationary-phase approximation, this lack of realism is not too important. To produce a realistic waveform at high frequencies, in a regime where the slow inspiral gives way to a rapid merger of the two stars, is still an open and challenging problem in gravitational-wave research.

ACKNOWLEDGMENTS

The authors express gratitude toward P. R. Brady, A. Gopakumar, B. S. Sathyaprakash, and A. G. Wiseman for discussions. The work in Guelph was supported by the Natural Sciences and Engineering Research Council of Canada. The work in Pasadena was supported by the National Science Foundation under the grant PHY-9424337, and under the NSF graduate program. E.P. wishes to thank Kip Thorne for his kind hospitality at the California Institute of Technology, where part of this work was carried out.

[1] K. S. Thorne, in *300 Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, England, 1987), p. 330.
[2] B. F. Schutz, in *The Detection of Gravitational Waves*, edited by D. G. Blair (Cambridge University Press, Cambridge, England, 1991), p. 406.
[3] B. S. Sathyaprakash and S. V. Dhurandhar, Phys. Rev. D **44**, 3819 (1991).
[4] S. V. Dhurandhar and B. S. Sathyaprakash, Phys. Rev. D **49**, 1707 (1994).
[5] C. Cutler and É. É. Flanagan, Phys. Rev. D **49**, 2658 (1994).
[6] R. Balasubramanian and S. V. Dhurandhar, Phys. Rev. D **50**, 6080 (1994).
[7] B. S. Sathyaprakash, Phys. Rev. D **50**, 7111 (1994).
[8] T. A. Apostolatos, Phys. Rev. D **52**, 605 (1996).
[9] E. Poisson and C. M. Will, Phys. Rev. D **52**, 848 (1995).
[10] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, Phys. Rev. D **53**, 3033 (1996).
[11] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).
[12] T. A. Apostolatos, Phys. Rev. D **54**, 2421 (1996).
[13] S. Droz and E. Poisson, Phys. Rev. D **56**, 4449 (1997).
[14] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, Phys. Rev. D **57**, 885 (1998).
[15] R. Balasubramanian and S. V. Dhurandhar, Phys. Rev. D **57**, 3408 (1998).
[16] B. J. Owen and B. S. Sathyaprakash, gr-qc/9808076.
[17] L. A. Wainstein and V. D. Zubakov, *Extraction of Signals from Noise* (Prentice-Hall, Englewood Cliffs, NJ, 1962).
[18] B. Allen *et al.*, GRASP software package, http://www.lsc-group.phys.uwm.edu/~ballen/grasp-distribution/.
[19] C. M. Bender and S. A. Orszag, *Advanced Mathematical Methods for Scientists and Engineers* (McGraw-Hill, New York, 1978), Chap. 6.
[20] L. Blanchet, B. R. Iyer, C. M. Will, and A. G. Wiseman, Class. Quantum Grav. **13**, 575 (1996).
[21] A. Gopakumar (in preparation).
[22] E. Chassande-Mottin and P. Flandrin, in IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis, 1998, p. 117.
[23] E. Chassande-Mottin and P. Flandrin, Appl. Comp. Harm. Anal. (to be published).
[24] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes* (Cambridge University Press, Cambridge, England, 1986), Chap. 12.