Mining information from binary black hole mergers: A comparison of estimation methods for complex exponentials in noise

Emanuele Berti[*]
McDonnell Center for the Space Sciences, Department of Physics, Washington University, St. Louis, Missouri 63130, USA

Vitor Cardoso[†]
Department of Physics and Astronomy, The University of Mississippi, University, Mississippi 38677-1848, USA

José A. González[‡] and Ulrich Sperhake[§]
Theoretical Physics Institute, University of Jena, Max-Wien-Platz 1, 07743, Jena, Germany
(Received 16 January 2007; published 21 June 2007)

The ringdown phase following a binary black hole merger is usually assumed to be well described by a linear superposition of complex exponentials (quasinormal modes). In the strong-field conditions typical of a binary black hole merger, nonlinear effects may produce mode coupling. Artificial mode coupling can also be induced by the black hole's rotation, if the radiation field is expanded in terms of spin-weighted spherical harmonics (rather than spin-weighted *spheroidal* harmonics). Observing deviations from the predictions of linear black hole perturbation theory requires optimal fitting techniques to extract ringdown parameters from numerical waveforms, which are inevitably affected by numerical error. So far, nonlinear least-squares fitting methods have been used as the standard workhorse to extract frequencies from ringdown waveforms. These methods are known *not* to be optimal for estimating parameters of complex exponentials. Furthermore, different fitting methods have different performance in the presence of noise. The main purpose of this paper is to introduce the gravitational wave community to modern variations of a linear parameter estimation technique first devised in 1795 by Prony: the Kumaresan-Tufts and matrix pencil methods. Using ''test'' damped sinusoidal signals in Gaussian white noise we illustrate the advantages of these methods, showing that they have variance and bias at least comparable to standard nonlinear least-squares techniques. Then we compare the performance of different methods on unequal-mass binary black hole merger waveforms. The methods we discuss should be useful both theoretically (to monitor errors and search for nonlinearities in numerical relativity simulations) and experimentally (for parameter estimation from ringdown signals after a gravitational wave detection).

DOI: 10.1103/PhysRevD.75.124017 PACS numbers: 04.25.Dm, 02.60.Ed, 04.30.Db, 04.70.Bw

I. INTRODUCTION

The past year witnessed a remarkable breakthrough in numerical relativity. Many different groups were finally able to evolve black hole binaries through the last few cycles of inspiral, merger, and ringdown and to extract gravitational waveforms from the evolutions [1–10]. If numerical simulations start out at large enough separation, it should be possible to match them with post-Newtonian predictions for the inspiral signal. This is a very active research area [11–13]. Studies of equal-mass merger simulations starting from different orbital separations (or produced by different numerical techniques) show that, independently of the waveform accuracy in the premerger phase, the strong-field waveform has a ''universal'' shape. Preliminary, ''first order'' explorations of the merger waveform for equal-mass black hole binaries indicate that the merger phase is short-lived, lasting only ~0.5–0.75 gravitational wave (GW) cycles [12]. After this short merger the signal has the typical ''ringdown'' shape—i.e., it is well modeled by a superposition of complex exponentials, known as quasinormal modes (QNMs).

QNMs play a role in any astrophysical process involving stars and black holes. Oscillations of a relativistic star and of a black hole necessarily produce GWs [14–16]. In linear perturbation theory the metric perturbations can usually be described by a single scalar function, which, in the ringdown phase, can be written as a superposition of complex exponentials:

$$\Psi_{lm}(t) = \sum_n A_{lmn} e^{(\alpha_{lmn} + i\omega_{lmn})t + i\varphi_{lmn}}. \qquad (1.1)$$

Here A_{lmn}, φ_{lmn}, ω_{lmn}, and $\tau_{lmn} = -\alpha_{lmn}^{-1}$ are the mode's amplitude, phase, frequency, and damping time, respectively. The indices (l, m) describe the angular dependence of the signal, and the index n sorts the modes by the magnitude of their damping time. The amplitude and phase is determined by the specific process exciting the oscillations. Remarkably, in linear perturbation theory the QNM frequencies ω_{lmn} and damping times τ_{lmn} are uniquely determined by the mass and angular momentum of the

[*]berti@wugrav.wustl.edu
[†]Also at Centro de Física Computacional, Universidade de Coimbra, P-3004-516 Coimbra, Portugal.
vcardoso@phy.olemiss.edu
[‡]jose.gonzalez@uni-jena.de
[§]Ulrich.Sperhake@uni-jena.de

black hole: this is a consequence of the so-called "no-hair" theorem of general relativity. For this reason, accurate QNM measurements could provide the "smoking gun" for black holes and an important test of general relativity in the strong-field regime (see e.g. [17], listing the dominant QNM frequencies for rotating black holes and providing fits of the numerical results).

In the GW literature it is often claimed that "the ringdown phase is well described by (linear) black hole perturbation theory, and that the ringdown waveform is given by a simple superposition of damped exponentials" of the form (1.1). This statement is based on two tacit assumptions, that may well be false under the strong-field conditions typical of a binary black hole merger:

(i) *Nonlinearities are negligible*, so that linear perturbation theory applies. However, nonlinear effects *should* be present in binary black hole merger waveforms. Model problems show that nonlinearities are responsible for a systematic shift in the quasinormal frequencies and, more interestingly, for mode coupling [18].

(ii) *Mode coupling is negligible*. Even if assumption (i) above is valid and nonlinear effects are small, different multipolar components can still be coupled. For example, artificial mode coupling can arise because of some commonly used approximations in the wave extraction process.[1]

An important open problem in the interpretation of merger waveforms is to isolate general relativistic nonlinearities from artificial effects induced by unwanted features of the numerical simulations, such as finite-differencing errors, errors due to the approximate nature of initial data, or spurious rotational QNM coupling induced by the use of spin-weighted spherical harmonics in GW extraction. All of these effects depend on the details of the numerical implementation and of the wave extraction procedure. The ultimate goal of GW detection is to assess the validity of general relativity in the strong-field regime, so the relevance of this problem can hardly be underestimated. Understanding nonlinear phenomena unique to general relativity is of paramount importance if we want to discriminate Einstein's theory from alternative theories of gravity, that may give different predictions in strong-field situations. Since most of the strong-field waveform can be described as a QNM superposition, strong-field effects could well show up in the fine structure of this QNM superposition: in particular, nonlinearities may manifest themselves as time variations of the ringdown frequencies, or as beating phenomena between different QNMs.

In this paper we take a first step towards the solution of this problem. We consider intrinsic limitations of the *fitting routines* used to extract quasinormal ringing parameters, exploring the properties of different fitting techniques for the extraction of ringdown parameters from numerical waveforms. This issue can be studied independently of the details of any given merger simulation. The idea is that, by quantifying the limitations of a given fitting method, we should be able to isolate *systematic uncertainties* (due to the fact that we are fitting noisy numerical waveforms) from possibly more interesting physical features of the waveforms themselves.

The problem of estimating the parameters of complex exponentials in noise has a long history in science and engineering. Damped sinusoidal signals are the "real-world counterpart" of the harmonic oscillators used in Fourier analysis. Therefore, estimating the parameters of damped sinusoids is of primary importance in physics. Examples include (i) speech and audio modeling [21], (ii) - angle-of-arrival estimation of plane waves impinging on sensor arrays for radar purposes or radio-astronomy, and antennae array processing [22], (iii) estimation of lifetimes and intensities in radioactive decay [23], (iv) nuclear magnetic resonance and computer assisted medical diagnosis [24]. This list can be extended to include spectroscopy, the study of mechanical vibrations (including seismic signal processing), the diffusion of chemical compounds, economics, and so on.

Fits of ringdown waveforms are usually performed using standard nonlinear least-squares methods (see e.g. [12,20] and references therein). These methods are known to fail in estimating parameters for a *sum* of damped exponentials. In this case, minimizing the squared error over the data requires the solution of highly nonlinear expressions in terms of sums of powers of the damping coefficient. No analytic solution is available, and these expressions can usually be solved only with a good initial guess for the parameters [25]. Even for a single damped exponential, if the initial guess is inaccurate the algorithm often fails to converge. Iterative algorithms, such as gradient descent procedures or Newton's method, have been devised to minimize these nonlinear expressions. Unfortunately, the algorithms are computationally very expensive, sometimes requiring at each step the inversion of matrices of dimension as large as the number of data samples [26]. Furthermore, gradient descent algorithms for multimodal equations sometimes fail to converge to the global minimum.

Computational difficulties with nonlinear least-squares methods led to the development of suboptimal estimation

[1]Spin-weighted spherical harmonics are often used to separate the angular dependence of the Weyl scalars and to read off the (l, m) multipolar components of the emitted radiation, but the Kerr metric is not spherically symmetric. More correctly, one should use spin-weighted *spheroidal* harmonics [19]. Spin-weighted spherical harmonics can be expressed as linear superpositions of spin-weighted spheroidal harmonics, and this leads to some artificial mode mixing in the extracted waveforms [12,20]. By "artificial" we mean that the effect—in this case, mode coupling—is due to approximations we introduce in the simulations, rather than being produced by a physical agent (such as nonlinearities).

methods based on linear prediction equations [24,25]. These methods are modern variations on a 1795 paper by Gaspard Riche, Baron de Prony [27], who introduced a procedure to *exactly* fit N data points by as many purely damped exponentials as needed. Modern versions of Prony's method generalize the original idea to damped sinusoidal models. They also make use of least-squares analysis to *approximately* fit an exponential model for cases where the data points cannot be fitted by the assumed number of exponential terms. Unfortunately, these "modified least-squares Prony methods" are very sensitive to numerical noise. Two successful improvements of these methods, the Kumaresan-Tufts (KT) [28] and matrix pencil (MP) [29] techniques, make use of singular value decomposition to improve parameter estimation accuracy in the presence of noise. For excellent reviews of these and other estimation methods we refer the reader to the first chapter of [24] (in French) and to Marple's book [25]. A purpose of this paper is to introduce these estimation methods to the GW community.

The plan of the paper is as follows. To put our problem in perspective, in Sec. II we describe the main features of our numerical simulations of binary black hole mergers and of the resulting waveforms. In Sec. III we summarize the theory behind different estimation methods, and in Sec. IV we list the numerical algorithms we implemented in our comparisons. In Sec. V we compare the performance of linear estimation methods (MP and KT) against nonlinear least-squares methods, by performing Monte Carlo simulations of damped sinusoidal signals in Gaussian white noise. In Sec. VI, we apply different fitting methods to selected waveforms generated by numerical relativity simulations. Finally, we summarize our results and indicate possible directions for further research.

II. NUMERICAL SIMULATIONS

In this section we briefly describe the numerical simulations of nonspinning, unequal-mass black hole binaries previously presented in [9]. The simulations were performed with the BAM code [8], using the so-called "moving puncture" method. This method, originally introduced in Refs. [3,4], is now being used almost routinely by various groups to successfully perform numerical simulations of inspiralling and merging black hole binaries. Our simulations of the inspiral of unequal-mass black hole binaries are the same used to study recoil in Ref. [9]. Details of these simulations, of the numerical setup, and of the BAM code are given in Refs. [8,9]. The main purpose of this paper is to test fitting methods to extract QNMs from the simulations. Therefore we defer a detailed investigation of the unequal-mass waveforms and of their physical properties to a forthcoming publication [30].

In the absence of spin, a standard conformally flat black hole binary initial data set is uniquely determined by the parameters m_1, m_2, $\vec{P}_1$, $\vec{P}_2$, D which denote, in this order, the bare masses and linear momenta of the individual holes, and their coordinate separation. The specification of these parameters allows us to compute the total Arnowitt-Deser-Misner (ADM) mass M of the system, the orbital angular momentum $L = (P_1 + P_2)D$, the individual black hole masses M_1, M_2, and the mass ratio $q = M_1/M_2$. Finally, the sum of the individual holes' spins and of the orbital angular momentum gives the total ADM angular momentum J of the system.

TABLE I. Total ADM mass, angular momentum, and initial binding energy of the unequal-mass binaries.

q	M	J/M^2	$-E_b/M$
1.00	0.9935	0.8845	0.0140
1.49	1.2422	0.8494	0.0134
1.99	1.4914	0.7870	0.0124
2.48	1.7408	0.7232	0.0114
2.97	1.9904	0.6649	0.0104
3.46	2.2401	0.6132	0.0096
3.95	2.4899	0.5679	0.0089

In the parameter study under consideration, the mass ratio q was varied while keeping constant the initial coordinate separation $D = 7M$ of the black holes. The linear momenta of the individual holes then follow from the requirement that the system represent a quasicircular configuration. For a given separation D, the momentum of each puncture can thus be calculated to third-order post-Newtonian accuracy in the Arnowitt-Deser-Misner, transverse-traceless (ADMTT) gauge. The resulting relation is given by Eq. (64) in [8] and forms the basis for all initial data sets discussed in this work. The parameters used in our simulations are summarized in Table I.

In the notation of Ref. [8], the simulations have been performed using the $\chi_{\eta=2}$ approach of the moving puncture method with a number of grid points $i = 56, 64, 72$ on the three innermost refinement levels, respectively. This corresponds to resolutions of $h = 1/45, 1/51$, and $1/58$. In this work we use the waveforms resulting from the low-resolution and high-resolution simulations, using $i = 56$ and $i = 72$ grid points, respectively. The wave extraction procedure is based on the Newman-Penrose formalism. The Weyl scalar Ψ_4 is extracted according to the method described in Sec. III A of [8]. In all simulations presented in this work, Ψ_4 is calculated at extraction radius $r_{ex} = 30M$. Further details of the numerical setup are given in Ref. [9].

To illustrate some issues related with the extraction of ringdown information from numerical waveforms, in Figs. 1 and 2 we show the complex mode amplitudes[2]

[2]These amplitudes are projections of the Weyl scalar Ψ_4 onto spin-weighted spherical harmonics of spin-weight $s = -2$. For the definition see, e.g., Eq. (39) of [8], where the amplitudes are denoted by A_{lm}. Here we choose a different notation, to avoid confusion with the QNM amplitudes A_{lmn}.

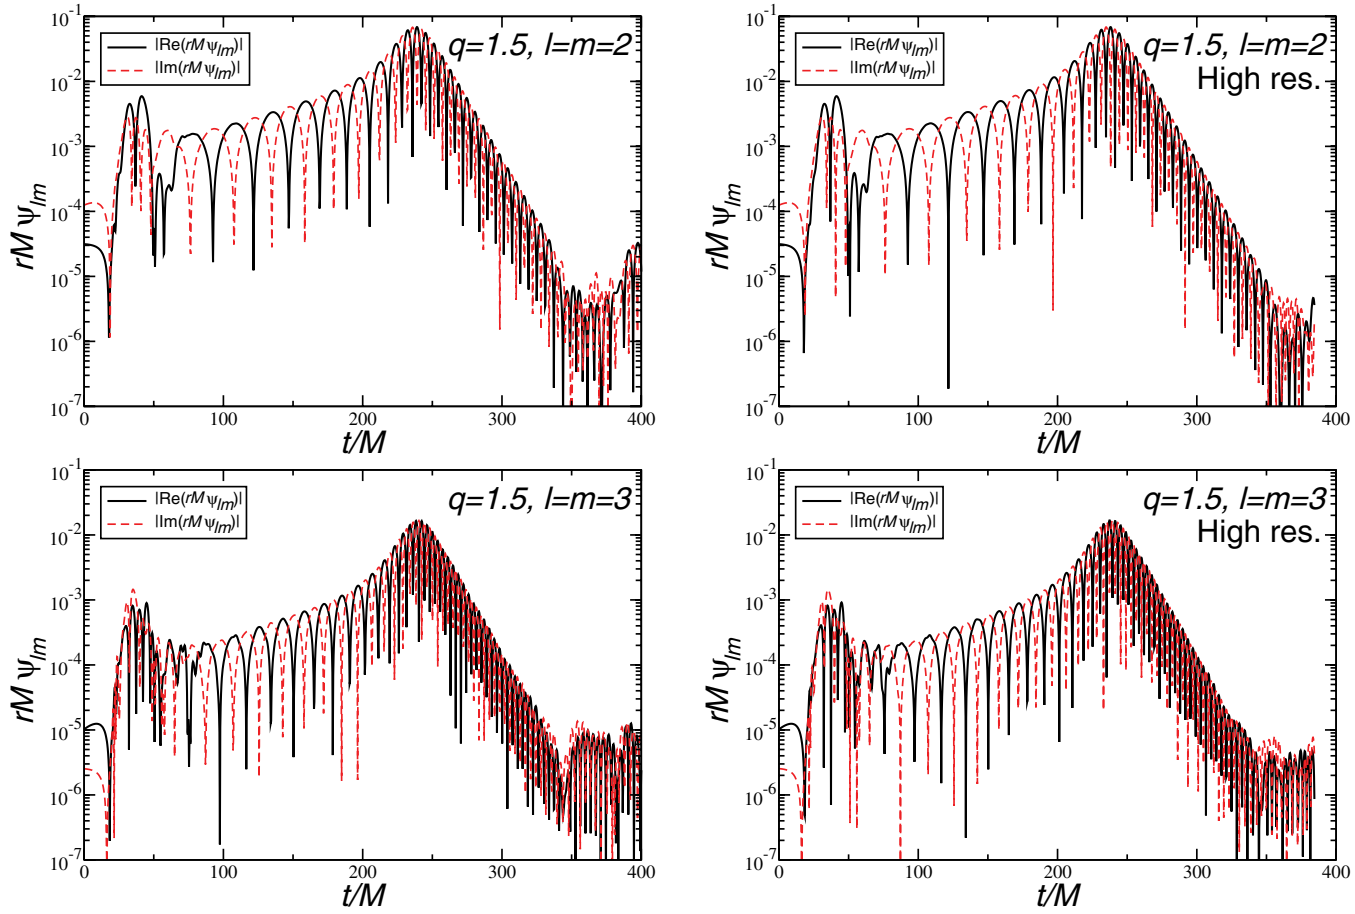

FIG. 1 (color online). Real and imaginary parts of $rM\psi_{lm}$ for the two dominant multipoles. The binary has mass ratio $q = 1.5$. The burst of radiation at early times is induced by the initial data. After the initial burst, the real and imaginary parts are simply related by a phase shift (i.e., the waveform is circularly polarized). The irregular behavior at late times is due to numerical noise.

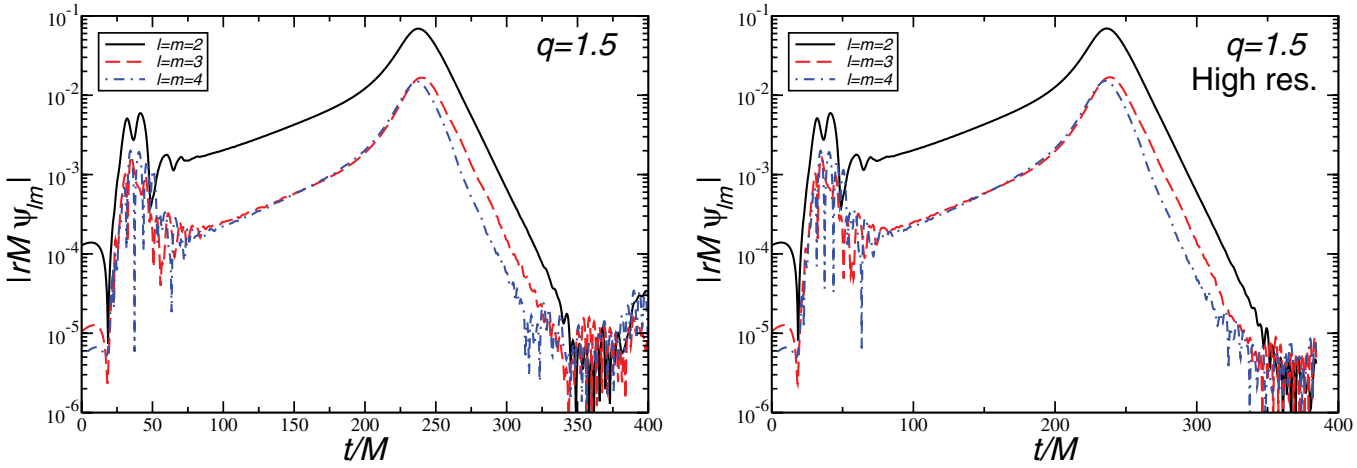

FIG. 2 (color online). $|rM\psi_{lm}|$ for different multipoles. The high-resolution $l = m = 2$, $l = m = 3$, and $l = m = 4$ wave amplitudes have maxima at $t_{\rm peak}/M = 236.4$, 238.4, and 235.7, respectively. Wiggles at late times are due to numerical noise, mainly caused by reflections from the boundaries.

$rM\psi_{lm}$ and their modulus $|rM\psi_{lm}|$ for the dominant components of the radiation emitted by a binary with mass ratio $q = 1.5$.[3] Because of the reflection symmetry of the system (see [8]), real and imaginary parts of the positive-m and negative-m components are related by

$$\psi_{l-m} = (-1)^l(\psi_{lm})^*. \tag{2.1}$$

Components with $|m| < l$ or $l > 4$ usually contain significant numerical noise, and we will ignore them in the following.

At early times the waveform is contaminated by a spurious burst of radiation, due to the approximate nature of the initial data. After this initial data burst, the frequency and amplitude of the wave grow as the binary members come closer, producing the characteristic "chirping" gravitational waveform. Eventually the binary members merge, the wave amplitude reaches a maximum, and then it decays as the remnant black hole settles down to a stationary Kerr state, emitting ringdown waves. The final part of the ringdown waveform is visibly contaminated by some amount of numerical noise, that decreases as we increase the resolution. This noise is mainly due to radiation being reflected from the boundaries of the computational domain.

The plots clearly show that the real and imaginary parts of the waveform follow the same pattern, except for a (roughly) constant phase shift. This means that the waveform is circularly polarized [31]. From the point of view of extracting information from the ringdown, circular polarization means that fitting the real or the imaginary part should make no difference: to a good approximation, we should get the same results for the oscillation frequencies and damping times. Fitting methods capable of directly dealing with *complex* waveforms (such as the Prony-type methods considered in this paper) should be particularly useful for waveforms with general polarization, such as those that should be emitted by the generic merger of spinning, precessing black holes.

Independently of the chosen fitting method, there is some arbitrariness in the choice of the time window $[t_0, t_f]$ used to perform the fit. A well-known problem with the transition merger-ringdown is that we do not know *a priori* when the ringdown starts (see [12,20,32] and references therein). Ideally, the starting time t_0 should be determined by a compromise between the following requirements: (i) t_0 should be small enough to include the largest possible number of data points: in particular, we do not want to miss the large amplitude, strong-field part of the waveform right after the merger; (ii) t_0 should be large enough that we do not include parts of the waveform which are *not* well described by a superposition of complex exponentials: the inclusion of inspiral and merger in the ringdown waveform would produce a bias in the QNM frequencies.

A judicious choice of t_f is also necessary. Usually we would like the time window to be as large as possible, but a glance at Fig. 1 and 2 shows that the low-amplitude, late-time signal is usually dominated by numerical noise, mainly caused by reflections from the boundaries. This noise can reduce the quality of the fit, especially for the subdominant components with $l > 2$ and for large values of t_0. A practical criterion for the choice of t_f is suggested by a look at Fig. 2. If the ringdown waveform were not affected by noise from boundary reflections, $|rM\psi_{lm}|$ should decay linearly on the logarithmic scale of the plots.[4] At low signal amplitudes, we see boundary noise-induced wiggles superimposed on this linear decay: the first occurrence of these wiggles is a good indicator of the time t_f at which numerical results cannot be trusted anymore. To test the robustness of fitting results to late-time numerical noise, while at the same time keeping the largest number of data points in the waveform, we decided to use two different "cutoff criteria":

(1) "Relative" cutoff: remove from the waveforms all data for times $t > t_f = t_{\rm rel}$, where $t_{\rm rel}$ is the time when the amplitude of each multipolar component $|rM\psi_{lm}|$ becomes less than some factor $\psi_{\rm cutoff}$ times the peak amplitude (at $t_{\rm peak} \sim 240M$ for the waveforms in Fig. 2):

$$\frac{|rM\psi_{lm}(t_{\rm rel})|}{|rM\psi_{lm}(t_{\rm peak})|} < \psi_{\rm cutoff}. \tag{2.2}$$

(2) "Absolute" cutoff: remove from the fit all data with $t > t_f = t_{\rm abs}$, where $t_{\rm abs}$ is the time at which the *absolute* value of the amplitude $|rM\psi_{lm}| < \psi_{\rm cutoff}/10$.

The choice of the cutoff amplitude is somewhat arbitrary. We chose $\psi_{\rm cutoff} = 10^{-3}$ for low resolution, and $\psi_{\rm cutoff} = 10^{-4}$ for high resolution.

For each chosen t_f, we will look at the performance of the fitting routines as we let t_0 vary in the range $[t_{\rm peak}, t_f]$. We do this for two reasons. The first reason is physical: by monitoring the convergence of the QNM frequencies to some "asymptotic" value as $t_0 \to \infty$, we can tell if the black hole settles down to a stationary Kerr state, or if, on the contrary, nonlinearities and mode coupling are always present. The second reason is related with the main goal of this paper, which is to assess the performance of our fitting

[3]Here and in the following we label the BAM runs by the corresponding value of q, rounded to the first decimal digit (see Table I).

[4]With larger resolution and longer running times, eventually the exponential decay should turn into the well-known power-law tail induced by backscattering of the radiation off the space-time curvature [33]. In the simulations we consider, noise produced by boundary effects is large enough that this effect is not visible.

routines: as t_0 grows the signal amplitude decreases exponentially, and we effectively reduce the signal-to-noise ratio (SNR) in our fitting window. Robust fitting methods should give reasonable results even for large values of t_0 (that is, modest values of the SNR).

III. A BRIEF SURVEY OF ESTIMATION METHODS FOR DAMPED SINUSOIDS

In this section we present a brief summary of the theory behind different estimation methods for damped sinusoidal signals. We consider only methods that make direct use of the data, by which we mean that no use is being made of the autocorrelation function. The measured signal x is a linear superposition of the "true" waveform, Ψ, and noise, ϖ. We suppose we have N samples of the signal, equally spaced in time with time sampling interval T, and we label each sample by an integer n:

$$x[n] = \Psi[n] + \varpi[n], \qquad n = 1, \ldots, N. \tag{3.1}$$

For simplicity we assume the noise $\varpi[n]$ to be white and Gaussian, with standard deviation σ and mean $\mu = 0$. We are interested in the ringdown waveform $\Psi[n]$, that we write as a superposition of p complex exponentials with arbitrary amplitudes and phases:

$$\Psi[n] = \sum_{k=1}^{p} A_k e^{(\alpha_k + i\omega_k)(n-1)T + i\varphi_k}. \tag{3.2}$$

It is useful to recast Eq. (3.2) in the slightly different form

$$\Psi[n] = \sum_{k=1}^{p} h_k z_k^{n-1}, \tag{3.3}$$

with

$$h_k = A_k e^{i\varphi_k}, \tag{3.4}$$

$$z_k = e^{(\alpha_k + i\omega_k)T}. \tag{3.5}$$

Now the unknown complex parameters are $\{h_k, z_k\}$ ($k = 1, \ldots, p$), and possibly the number p of damped sinusoids.

A. Nonlinear least squares

A popular estimation method is nonlinear least squares, which consists in minimizing the integrated squared error

$$\rho \equiv \sum_{n=1}^{N} |\varpi[n]|^2, \tag{3.6}$$

with $\varpi[n] = x[n] - \Psi[n]$. The method is very general, in the sense that (in principle) it can be applied to any model function $\Psi[n]$. For ringdown waveforms of the form (3.3) we must minimize over the $\{h_k, z_k\}$ parameter space (and possibly the p-space, if the number of damped exponentials is not known *a priori*) in an essentially nontrivial way. The procedure is computationally expensive and not always accurate: the solution may converge to a local (rather than a global) minimum of the integrated squared error. Experiments show that nonlinear least squares techniques perform badly in estimating the parameters of damped exponentials. The situation gets even worse in the presence of noise [24,25]. In our numerical work, to minimize the integrated squared error we used the well-known Levenberg-Marquardt algorithm [34] as implemented in the FORTRAN subroutine `lmdif`, which is part of the `MINPACK` library for solving systems of nonlinear equations [35].[5]

B. Prony method

Computational difficulties with nonlinear least-squares methods led to the development of suboptimal estimation methods, especially designed to deal with damped sinusoidal signals. The prototype of these estimation techniques is the Prony method, which is essentially a trick to reduce the nonlinear minimization problem to a linear prediction problem. The method has been successfully tested in different branches of data analysis and signal processing. Some variants of the basic idea improve the variance and bias of parameter estimation in the presence of noise: we briefly discuss these variants in the following sections. Our introduction to standard Prony methods parallels those by Marple [25] and Djermoune [24]. We refer the reader to those references for further details.

To start with, let us assume there is no noise in our data set. Let us also assume that there are as many data samples as there are complex exponential parameters $(h_1, \ldots, h_p)$, $(z_1, \ldots, z_p)$: $N = 2p$. We are then fitting $x[n]$ to an exponential model,

$$x[n] = \sum_{k=1}^{p} h_k z_k^{n-1}. \tag{3.7}$$

For $1 \le n \le p$ we can write this in matrix form as

$$\begin{pmatrix} z_1^0 & z_2^0 & \cdots & z_p^0 \\ z_1^1 & z_2^1 & \cdots & z_p^1 \\ \vdots & \vdots & \ddots & \vdots \\ z_1^{p-1} & z_2^{p-1} & \cdots & z_p^{p-1} \end{pmatrix} \begin{pmatrix} h_1 \\ h_2 \\ \vdots \\ h_p \end{pmatrix} = \begin{pmatrix} x[1] \\ x[2] \\ \vdots \\ x[p] \end{pmatrix}. \tag{3.8}$$

If we can determine the z_k's by some other procedure, then (3.8) is a set of linear equations for the complex amplitudes h_k. Prony's method is in essence a trick to determine the z_k's without the need for nonlinear minimizations, as follows. Let us define a polynomial $\mathbf{A}(z)$ of degree p which has the z_k's as its roots:

$$\mathbf{A}(z) = \prod_{k=1}^{p} (z - z_k) \equiv \sum_{m=0}^{p} a[m] z^{p-m}. \tag{3.9}$$

Let us also normalize the complex coefficients $a[m]$ so that

[5]The same algorithm is also implemented in MATHEMATICA [36].

$a[0] = 1$. By using Eq. (3.7) we can write

$$a[m]x[n-m] = a[m]\sum_{k=1}^{p} h_k z_k^{n-m-1}. \tag{3.10}$$

Summing this last equation from $m = 0$ to $m = p$ we get, for $p+1 \le n \le 2p$

$$\sum_{m=0}^{p} a[m]x[n-m] = \sum_{i=1}^{p} h_i z_i^{n-p} \sum_{m=0}^{p} a[m] z_i^{p-m-1} = 0, \tag{3.11}$$

where the last equality follows from (3.9). This is a *(forward) linear prediction equation*. In matrix form it can be expressed as

$$\begin{pmatrix} x[p] & x[p-1] & \cdots & x[1] \\ x[p+1] & x[p] & \cdots & x[2] \\ \vdots & \vdots & \ddots & \vdots \\ x[2p-1] & x[2p-2] & \cdots & x[p] \end{pmatrix} \begin{pmatrix} a[1] \\ a[2] \\ \vdots \\ a[p] \end{pmatrix} = -\begin{pmatrix} x[p+1] \\ x[p+2] \\ \vdots \\ x[2p] \end{pmatrix}. \tag{3.12}$$

This equation is the basic result of the Prony method: it shows that we can determine the $a[k]$'s from the data, decoupling the problem of determining the h_k and z_k parameters.

The original Prony procedure works as follows. First, given the data, solve (3.12) for the coefficients $a[m]$. Then determine the roots z_k of the polynomial (3.9). The damping and frequency are obtained by inverting (3.5):

$$\alpha_k = \log|z_k|/T, \tag{3.13a}$$
$$\omega_k = \tan^{-1}[\mathrm{Im}(z_k)/\mathrm{Re}(z_k)]/T. \tag{3.13b}$$

Finally, solve (3.8) for the complex quantities h_k and invert (3.4) to find amplitudes and phases:

$$A_k = |h_k|, \tag{3.14a}$$
$$\varphi_k = \tan^{-1}[\mathrm{Im}(h_k)/\mathrm{Re}(h_k)]. \tag{3.14b}$$

It should now be clear that the Prony algorithm reduces the nonlinear fitting problem to the trivial, computationally inexpensive numerical tasks of (i) solving linear systems of equations, and (ii) finding the roots of a polynomial.

C. Modified least-squares Prony

For most situations of interest, there are more data points than there are exponential parameters: $N > 2p$. In this case, Eq. (3.12) is modified to

$$\begin{pmatrix} x[p] & x[p-1] & \cdots & x[1] \\ x[p+1] & x[p] & \cdots & x[2] \\ \vdots & \vdots & \ddots & \vdots \\ x[N-1] & x[N-2] & \cdots & x[N-p] \end{pmatrix} \begin{pmatrix} a[1] \\ a[2] \\ \vdots \\ a[p] \end{pmatrix} = -\begin{pmatrix} x[p+1] \\ x[p+2] \\ \vdots \\ x[N] \end{pmatrix}. \tag{3.15}$$

We can write this as a matrix equation,

$$\mathbf{X}\,\mathbf{a} = -\mathbf{x}, \tag{3.16}$$

which can be solved in the least-squares sense

$$\mathbf{a} = -(\mathbf{X}^H\mathbf{X})^{-1}\mathbf{X}^H\mathbf{x}, \tag{3.17}$$

where a superscript H denotes Hermitian transpose: $\mathbf{X}^H = (\mathbf{X}^T)^*$. The important difference is that we must now minimize the *linear prediction* squared error, rather than the *exponential approximation* squared error of Eq. (3.6). Once we have determined the $a[k]$'s, the rest of the Prony method carries on in the same manner. Details of the implementation of this "least-squares Prony method" (which is sometimes called "extended Prony method") can be found in [25].

D. Kumaresan-Tufts

Unfortunately, the original and least-squares Prony methods only work well in the absence of noise, or for large SNR. For small SNR, the estimation of the frequencies and damping times has large variance and bias [25,37].

In the previous section we explained how to estimate the exponential parameters by introducing the characteristic polynomial $\mathbf{A}(z)$, which has roots at $z_k \equiv e^{s_k} = e^{(\alpha_k + i\omega_k)T}$. The coefficients $a[k]$ of $\mathbf{A}(z)$ are solutions of the forward linear prediction equation (3.11). These same exponentials can be generated in reverse time by the backward linear predictor [25]

$$\sum_{m=0}^{p} b[m]x[n-p+m] = 0. \tag{3.18}$$

in which $b[0] = 1$. The characteristic polynomial

$$\mathbf{B}(z) = \sum_{m=0}^{p} b^*[m] z^{p-m} \tag{3.19}$$

has roots at $z_k = e^{-s_k^*}$. For damped sinusoids ($\mathrm{Re}[\alpha_k] < 0$) it can be shown that the roots of the forward linear prediction polynomial $\mathbf{A}(z)$ lie *inside* the unit z-plane circle, whereas those of the backward linear prediction polynomial $\mathbf{B}(z)$ lie *outside* the unit z-plane circle [25,28].

Suppose now we superimpose to the signal complex additive Gaussian white noise. Noise causes a bias in the estimates of the true zeros of the polynomials, which translates into a bias in the estimates of frequencies and damping times. It was empirically observed that the bias may be significantly reduced by looking for a number of exponential components $L > p$, where p is the actual number of exponentials present in the signal [28]. L is called the *prediction order* of the model. Of course, the selection of a higher prediction order introduces additional zeros due to noise, but these can be statistically separated by examining the zeros of $\mathbf{A}(z)$ and $\mathbf{B}(z)$. For both polynomials, zeros due to the noise tend to stay within the unit circle, whereas the true zeros due to the exponential signal form complex-conjugate pairs inside and outside the unit circle. This is basically due to the fact that the statistics of a stationary random process do not change under time reversal. Using singular value decomposition (SVD) can provide further improvement [24,25,28]. Write $\mathbf{X}$ in (3.16) as

$$\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^H, \tag{3.20}$$

where $\mathbf{S}$ is a $(N-L)\times L$ dimensional matrix with the singular values on the diagonal $(s_1, \ldots, s_p, s_{p+1}, \ldots, s_L)$ arranged in decreasing order. Noise can be reduced by considering the reduced rank approximation

$$\hat{\mathbf{X}} = \mathbf{U}\hat{\mathbf{S}}\mathbf{V}^H \tag{3.21}$$

with

$$\hat{\mathbf{S}} = \begin{bmatrix} \mathbf{S}_p & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}_{(N-L)\times L}. \tag{3.22}$$

Here $\mathbf{S}_p$ is the top-left $p \times p$ minor of $\mathbf{S}$. An estimate for the coefficients $a[k]$ is then

$$\hat{\mathbf{a}} = -\hat{\mathbf{X}}^{+}\mathbf{x}, \tag{3.23}$$

where $\hat{\mathbf{X}}^{+}$ is the pseudoinverse of $\hat{\mathbf{X}}$. It was found [28] that the use of the truncated SVD greatly enhances the SNR, providing a better estimate of the vector $\hat{\mathbf{a}}$, and consequently of the exponential parameters. This is the basic idea underlying the Kumaresan-Tufts method. More details can be found in the original work [28] (see also [24,25]). For improvements of this method using total least squares, see for instance [38].

The frequency variance depends on the prediction order L. For one *undamped* exponential ($\alpha = 0$) of amplitude h_1 it can be shown [29,39] that the variance is given by

$$\mathrm{var}(\omega) = \frac{2(2L+1)}{3(N-L)^2L(L+1)}\frac{\sigma^2}{|h_1|^2} \quad \text{for } L \le N/2, \tag{3.24a}$$

$$\mathrm{var}(\omega) = \frac{2[-(N-L)^2+3L^2+3L+1]}{3(N-L)L^2(L+1)^2}\frac{\sigma^2}{|h_1|^2} \quad \text{for } L \ge N/2. \tag{3.24b}$$

Minima are attained for $L \simeq N/3$ ($L \le N/2$) and for $L \simeq 2N/3$ ($L \ge N/2$), these relations being exact in the limit $N \to \infty$. Correspondingly, the optimal frequency variance is

$$\mathrm{var}\,(\omega) \simeq \frac{27}{4N^3}\frac{\sigma^2}{|h_1|^2}, \tag{3.25}$$

which is only slightly larger than the Cramer-Rao bound:

$$\mathrm{CRB}\,(\omega) = \frac{6}{N(N^2-1)}\frac{\sigma^2}{|h_1|^2}, \tag{3.26}$$

the ratio being $9/8 = 1.125$ in the limit $N \to \infty$. For one damped exponential, closed-form expressions for the variance and bias of the damping time in the large SNR limit can be found in [40,41].

E. Matrix pencil

An alternative to estimate exponential parameters from noisy signals is the matrix pencil (MP) method [29]. The MP method is in general more robust than the KT method, having a lower variance on the estimated parameters, but a slightly larger bias [24,41]. For a detailed description of this method we refer the reader to [29]. The following brief summary follows quite closely the treatment in [24]. Let $\mathbf{X}_0$ and $\mathbf{X}_1$ be two matrices defined as

$$\mathbf{X}_0 = \begin{pmatrix} x[0] & x[1] & \cdots & x[L-1] \\ x[1] & x[2] & \cdots & x[L] \\ \vdots & \vdots & \ddots & \vdots \\ x[N-L-1] & x[N-L] & \cdots & x[N-2] \end{pmatrix};$$

$$\mathbf{X}_1 = \begin{pmatrix} x[1] & x[2] & \cdots & x[L] \\ x[2] & x[3] & \cdots & x[L+1] \\ \vdots & \vdots & \ddots & \vdots \\ x[N-L] & x[N-L] & \cdots & x[N-1] \end{pmatrix}. \tag{3.27}$$

Here L is the pencil parameter: it plays the role of the prediction order parameter in the KT method. One can decompose $\mathbf{X}_0$ and $\mathbf{X}_1$ as

$$\mathbf{X}_0 = \mathbf{Z}_l\mathbf{H}\mathbf{Z}_r, \tag{3.28}$$

$$\mathbf{X}_1 = \mathbf{Z}_l\mathbf{H}\mathbf{Z}\mathbf{Z}_r, \tag{3.29}$$

where

$$\mathbf{Z}_l = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ z_1 & z_2 & \cdots & z_p \\ \vdots & \vdots & \ddots & \vdots \\ z_1^{N-L-1} & z_2^{N-L-1} & \cdots & z_p^{N-L-1} \end{pmatrix}; \qquad \mathbf{Z}_r = \begin{pmatrix} 1 & z_1 & \cdots & z_1^{L-1} \\ 1 & z_2 & \cdots & z_2^{L-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & z_p & \cdots & z_p^{L-1} \end{pmatrix}, \tag{3.30}$$

$$\mathbf{H} = \mathrm{diag}(h_1, h_2, \ldots, h_p), \tag{3.31}$$

$$\mathbf{Z} = \mathrm{diag}(z_1, z_2, \ldots, z_p). \tag{3.32}$$

Consider now the matrix pencil[6] $\mathbf{X}_1 - z\mathbf{X}_0$. We can write this as

$$\mathbf{X}_1 - z\mathbf{X}_0 = \mathbf{Z}_l\mathbf{H}(\mathbf{Z} - z\mathbf{I}_p)\mathbf{Z}_r. \tag{3.33}$$

When $z \neq z_i$, the matrix $\mathbf{Z} - z\mathbf{I}_p$ is of rank p. However, for $z = z_i$ it is of rank $p-1$. Therefore the poles of the signal reduce the rank of the matrix pencil for $p \leq L \leq N - p$. This is equivalent to saying that the poles z_i are the generalized eigenvalues of $(\mathbf{X}_1, \mathbf{X}_0)$, in the sense that $(\mathbf{X}_1 - z\mathbf{X}_0)\mathbf{v} = 0$, with $\mathbf{v}$ an eigenvector of $\mathbf{X}_1 - z\mathbf{X}_0$. To find the poles z_i one can use the fact that $\mathbf{X}_0^+\mathbf{X}_1$ has p eigenvalues equal to the poles z_i, and $L - p$ null eigenvalues [24]. Here a dagger denotes the (Moore-Penrose) pseudoinverse.

In practice we do not have access to the noiseless signal, therefore we must work directly with the noisy data. SVD is again necessary to select the singular values due to the signal. The basic steps of the MP method can be summarized as follows:

(i) Build the matrices $\mathbf{Y}_0$ and $\mathbf{Y}_1$ as in (3.27).
(ii) Make a SVD of $\mathbf{Y}_1$: $\mathbf{Y}_1 = \mathbf{U}\mathbf{S}\mathbf{V}^T$.
(iii) Estimate the signal subspace of $\mathbf{Y}_1$ by considering the p largest singular values of $\mathbf{S}$: $\tilde{\mathbf{Y}}_1 = \mathbf{U}_p\mathbf{S}_p\mathbf{V}_p^H$, where $\mathbf{U}_p$ and $\mathbf{V}_p$ are built from the first p columns of $\mathbf{U}$ and $\mathbf{V}$, and $\mathbf{S}_p$ is the top-left $p \times p$ minor of $\mathbf{S}$.
(iv) The matrix $\mathbf{Z}_L = \mathbf{Y}_1^+\mathbf{Y}_0 = \mathbf{V}_p\mathbf{S}_p^{-1}\mathbf{U}_p^T\mathbf{Y}_0$ has p eigenvalues which provide estimates of the inverse poles $1/z_i$; the other $L - p$ eigenvalues are zero. Since $\mathbf{Z}_L$ has only p nonzero eigenvalues, it is convenient to restrict attention to a $p \times p$ matrix $\mathbf{Z}_p = \mathbf{S}_p^{-1}\mathbf{U}_p^T\mathbf{Y}_0\mathbf{V}_p$ [29].

The MP technique exploits the matrix pencil structure of the underlying signal, rather than the prediction equations satisfied by it. Nevertheless, there are strong similarities between the MP and KT methods. An extensive comparison of their performance and theoretical properties can be found in [24,29,41].

As for the KT method, the variance of parameter estimation depends on L. For one *undamped* exponential of amplitude h_1 it can be shown [29] that the variance is given by

$$\mathrm{var}(\omega) = \frac{1}{(N-L)^2 L}\frac{\sigma^2}{|h_1|^2} \quad \text{for } L \leq N/2, \tag{3.34a}$$

$$\mathrm{var}(\omega) = \frac{1}{(N-L)L^2}\frac{\sigma^2}{|h_1|^2} \quad \text{for } L \geq N/2, \tag{3.34b}$$

i.e. it is always *lower* than the corresponding variance for the KT method, as given in Eq. (3.24). Choices minimizing the frequency variance are $L = N/3$ for $L \leq N/2$, and $L = 2N/3$ for $L \geq N/2$, and the optimal frequency variance is

$$\mathrm{var}\,(\omega) = \frac{27}{4N^3}\frac{\sigma^2}{|h_1|^2}, \tag{3.35}$$

as in the KT method. For one damped exponential, closed-form expressions for the variance and bias of the damping time in the large SNR limit have recently been derived and compared with Monte Carlo simulations [41]. The results indicate that the MP method performs *better* (has smaller variance) than the KT method. For both methods the frequency estimate is unbiased, but the estimate of the damping time is biased. The bias is slightly larger for the MP method, the difference between the two methods becoming larger for small SNR.

IV. FITTING ALGORITHMS CONSIDERED IN THIS PAPER

In the rest of this paper we will compare the performance of the four fitting algorithms described in the previous section: Levenberg-Marquardt (LM), modified least-squares Prony, Kumaresan-Tufts (KT), and matrix pencil (MP). For quick reference, in Table II we list the original papers, along with web resources and publications providing software implementations of each algorithm (in

TABLE II. Fitting methods used in this paper. El-Hadi Djermoune and Stanley Lawrence Marple, Jr. kindly sent us up-to-date implementations of the codes described in [24,25]. We are also grateful to Gordon Smyth for providing us with a MATLAB routine to estimate parameters for purely damped exponentials, and to Boaz Porat for MATHEMATICA packages to estimate parameters of undamped sinusoids using several methods (KT, maximum likelihood, Yule-Walker ...): see [42].

Method	Reference	Software
Levenberg-Marquardt (LM)	[34]	[35,36]
Modified least-squares Prony	[25]	[25,43–45]
Kumaresan-Tufts (KT)	[28]	[24,46]
Matrix pencil (MP)	[29]	[24]

[6] Let A and B be two $n \times n$ matrices. The set of all matrices of the form $A - \lambda B$, with λ complex, is called a *matrix pencil*.

FORTRAN, MATLAB, or MATHEMATICA). Below we give some details on our own practical implementation of the algorithms.

A. Levenberg-Marquardt

The LM algorithm is more general than the other methods we consider in this paper, in the sense that fitting functions are not restricted to a simple sum of complex exponentials. A problem if we want to fit merger waveforms is that LM (like most nonlinear least-squares methods) is designed to deal with *real* functions. For this reason the authors of [12], who studied waveforms similar to those in Fig. 1, fitted only the real (or imaginary) part of the signal. Some practice also shows that general nonlinear least-squares methods (such as the LM algorithm) often fail to converge if the signal is given by a superposition of damped sinusoids, unless we provide very accurate initial guesses for the fitting parameters.

For simplicity, and to take into account these limitations of nonlinear least-squares algorithms, in Sec. V we will compare the performance of different routines on a *real* ringdown signal with a *single* frequency and damping time. We choose some starting time t_0 and we fit the real signal by a four-parameter function:

$$\Psi(t) = Ae^{\alpha(t-t_0)}\sin(\omega t + \varphi). \tag{4.1}$$

The LM algorithm provides reasonable answers only if the four parameters $(A, \varphi, \omega, \alpha)$ are reasonably close to their true values. In particular, an accurate initial guess for A is necessary to avoid ''hard failures'' of the fitting routine: we choose as an initial guess the value of the signal amplitude at t_0. We fix the tolerance parameter in the `lmdif` routine to be `TOL` $= 10^{-12}$ (but this choice is not crucial).

As a final remark we point out that, by using a real signal, in a sense we are ''biasing'' our tests in favor of the LM method. The reason is that fitting a real signal requires the inclusion of at least *two* complex exponentials in the sum (3.2), unnecessarily doubling the number of unknown parameters to be searched for by Prony methods.

B. Modified least-squares Prony, Kumaresan-Tufts, and matrix pencil

A FORTRAN routine implementing modified least-squares Prony was kindly provided to us by Stanley Lawrence Marple, Jr. [25]. We tested the routine's performance on noiseless ''pure ringdown'' waveforms obtained by superposing three damped exponentials with frequencies given by the first three QNMs ($n = 0, 1, 2$) of a Schwarzschild black hole [17]. For noiseless waveforms, the frequency and damping time of the fundamental mode ($n = 0$) are typically determined with accuracies $\sim 10^{-6}$. The frequency and damping time of the two overtones usually have accuracies $\sim 10^{-5}$. Parameter estimation becomes much worse, as expected from the theory, when we add even a modest amount of noise to the waveforms.

Performance in the presence of noise gets much better when we use two MATLAB routines by El-Hadi Djermoune [24], implementing the Kumaresan-Tufts and matrix pencil methods. These routines require the specification of the prediction order (that we always choose to be $L = N/3$, where N is the number of data points, in order to minimize the variance in parameter estimation). In addition, they also require the specification of the number of complex exponentials to be searched for. Djermoune kindly provided us with an additional routine to *estimate* the number of complex exponentials present in the signal, or (in more technical jargon) the ''order of the exponential model.'' The estimation routine is based on two criteria: Akaike's information criterion (AIC) and the minimum description length (MDL) criterion [47,48].

When tested on noiseless waveforms, Djermoune's routines are remarkably successful at estimating the number of modes present in the signal. Frequencies and damping times for noiseless signals are typically determined to accuracies $\sim 10^{-15}$ (comparable to machine precision).

We will see in the following that MP and KT methods have essentially the same performance. In fact, in the next section we will show by numerical experiments that MP methods have slightly smaller variance and bias in parameter estimation: this was pointed out also by Djermoune and collaborators (see e.g. [24,41,49]). Because of the similarity of the two techniques, and due to the slightly superior performance of MP over KT, we will usually compare MP and LM methods in our analysis of merger waveforms (Sec. VI below).

As a final remark, we point out that MP and KT have a variance in frequency estimation which is only $\simeq 9/8 = 1.125$ times larger than the Cramer-Rao bound. These algorithms may prove extremely useful not only to study the merger waveforms produced by numerical relativity (as we do here), but also to estimate the source parameters after GW detection. MP and KT algorithms are most effective and useful when we deal with noisy waveforms, and these situations are often the most interesting. For example, fitting the late-time, low-amplitude portion of a numerical relativity waveform yields the remnant black hole's parameters. Fitting large-l multipolar components (that carry little energy, hence are more affected by numerical noise) is necessary for tests of the general relativistic no-hair theorem [17]. Last but not least, the importance of parameter estimation from experimental, noisy GW data can hardly be underestimated.

V. COMPARISON OF DIFFERENT METHODS ON NOISY DAMPED SINUSOIDS

To test the performance of different fitting methods, we built model waveforms reproducing the essential features of the ringdown waveforms produced by binary black hole merger simulations (see Fig. 1). For simplicity, and to take into account the limitations of the LM algorithm, here we

consider our true signal as consisting of a single damped sinusoid. We take the frequency and damping time to be those of the fundamental ($n = 0$) $l = m = 2$ perturbations of a Kerr black hole, as listed in [17]. For a given oscillation frequency ω, we produce a signal lasting five GW cycles: i.e. the signal length $t_{\rm fin} = 5(2\pi/\omega) = 5T_{\rm GW}$. The sampling time is taken to be $T = T_{\rm GW}/50$, so that the "full" waveform consists of 250 data points.

To this true signal we superimpose Gaussian white noise. In this section we do not consider standard least-squares Prony, since it is outperformed by MP and KT algorithms in the presence of noise. The MP and KT algorithms are implemented in MATLAB. In this case, we produce Gaussian noise using the built-in function `normrnd`. For our FORTRAN implementation of the LM algorithm we generate noise as a random variable ϖ with mean $\mu = 0$ and standard deviation σ using the Box-Müller method [50]: for each t, given two random numbers u_1 and u_2 uniformly distributed in $[0, 1]$ (as generated with the Numerical Recipes routine `ran2` [51]), we add to our signal

$$\varpi(t) = \mu + \sigma\sqrt{-2\ln(u_1)}\cos(2\pi u_2). \tag{5.1}$$

For illustration, in Fig. 3 we show a typical waveform generated in this way. The true signal (with $\sigma = 0$) corresponds to a perturbed Schwarzschild black hole, i.e. the frequency and damping time correspond to the fundamental $l = m = 2$ QNM for a black hole of dimensionless spin parameter $j = 0$ [17]. For any given σ, the addition of Gaussian white noise has a less severe impact for large values of j. The reason is that the damping time of the $l = m = 2$ mode grows with the rotation parameter j (see e.g. Table II in [17]). A Schwarzschild waveform with $\sigma = 10^{-3}$ is quite similar to the typical merger waveforms of Fig. 1, and for this reason we will use it to test the performance of the different fitting routines. Results for different values of σ are qualitatively similar.

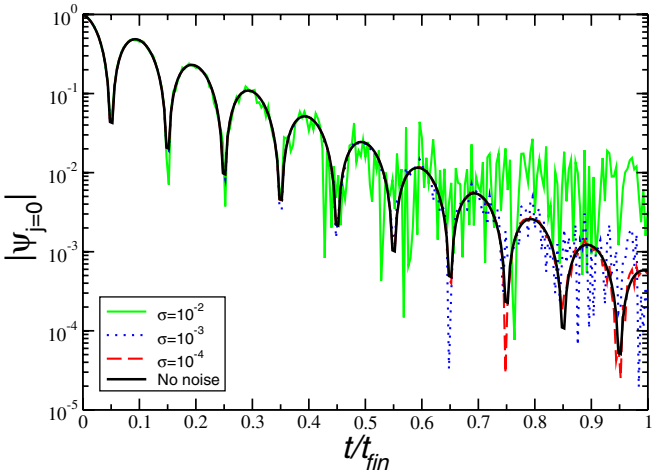

FIG. 3 (color online). QNM waveforms with and without Gaussian white noise. We normalize the time axis to the total duration of the signal $t_{\rm fin}$. The "pure" noiseless waveform (thick black line) has unit amplitude. Dashed (red), dotted (blue), and thin (green) lines are the same waveform superimposed to Gaussian white noise with $\sigma = 10^{-4}$, 10^{-3}, 10^{-2}, respectively.

For each fitting algorithm, we perform Monte Carlo simulations to produce a large number ($N_{\rm noise} = 100$) of realizations of the noise with $\sigma = 10^{-3}$. For each realization of the noise ($\nu = 1, \ldots, N_{\rm noise}$) we fit the waveform varying the starting time from $t_0 = 0$ to $t_0 = t_{\rm fin}$. Each fit yields a couple of values $\omega_\nu(t_0)$, $\alpha_\nu(t_0)$, and from these we can deduce the quality factor of the oscillation $Q_\nu(t_0) \equiv |\omega_\nu(t_0)/[2\alpha_\nu(t_0)]|$ [17]. For each t_0 we compute the standard deviation and bias of each of these fitted quantities x ($= \omega$, α or Q) from our Monte Carlo simulations:

$$\mathrm{std}\,(x)|_{t_0} = \left\{\frac{1}{N_{\rm noise} - 1}\sum_{\nu=1}^{N_{\rm noise}}[x_\nu(t_0) - \bar{x}(t_0)]^2\right\}^{1/2}, \tag{5.2}$$

$$\mathrm{bias}\,(x)|_{t_0} = \bar{x}(t_0) - x_{\rm true}, \tag{5.3}$$

where $x_{\rm true}$ is the known, true value of each quantity and a bar denotes the average over all noise realizations:

$$\bar{x}(t_0) = \left[\frac{1}{N_{\rm noise}}\sum_{\nu=1}^{N_{\rm noise}} x_\nu(t_0)\right]. \tag{5.4}$$

Results of these Monte Carlo simulations are shown in Fig. 4. There we see that all three methods (KT, MP, and LM) are essentially equivalent in terms of variance.[7] Standard deviation and bias grow sharply for $t_0/t_{\rm fin} \gtrsim 0.7$, when we are fitting that part of the signal which is buried in noise (compare Fig. 3). The bias is usually very small, but remarkably linear methods (KT, and especially MP) beat the nonlinear LM fitting routine in terms of bias on the frequency. In terms of bias on the damping time, LM performs slightly better than MP and KT only for large values of t_0, when the SNR is very small.

When comparing the performance of the different methods it is useful to remember that LM only works if we give a good initial guess for the parameters we want to estimate, whereas MP and KT *automatically* find the values of these parameters, with no need for initial guesses. Our comparison is somehow "biased" in favor of the LM method in at least three ways: (i) we choose optimal initial guesses for the parameters, so that the LM algorithm is not "allowed" to converge to some wrong root; (ii) we choose a real signal rather than a complex signal, because the LM algorithm only works for real data sets; (iii) we only consider one damped exponential (LM often fails to converge if we include additional damped sinusoids, whereas Prony-type methods are still remarkably successful).

[7]We also tested the MP and KT methods on pure sinusoids, finding good agreement with theoretical predictions for the frequency variance.

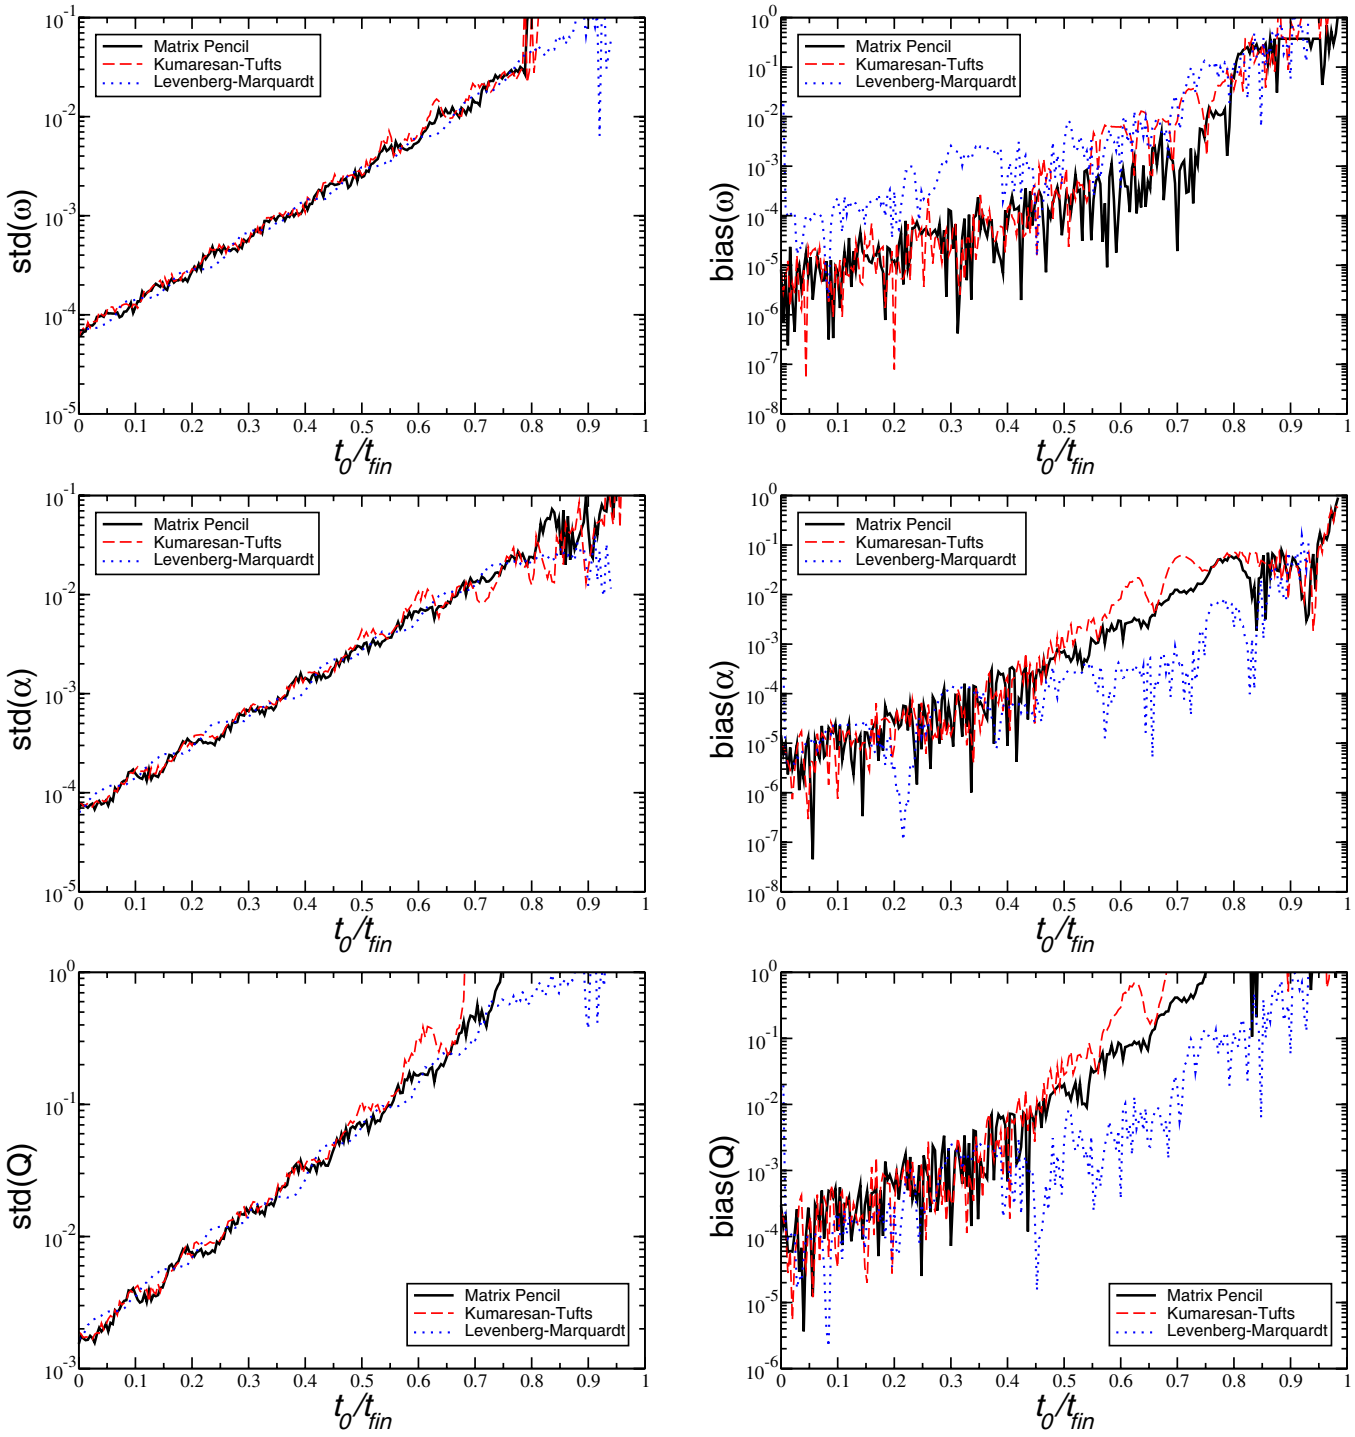

FIG. 4 (color online). Standard deviation (left) and bias (right) in the estimate of frequency, damping time, and quality factor (top to bottom). All quantities are given as functions of the starting time of the fit t_0 (normalized by the duration of the signal $t_{\rm fin}$); each point is the result of a Monte Carlo simulation obtained by adding $N_{\rm noise} = 100$ realizations of Gaussian white noise with zero mean and $\sigma = 10^{-3}$ to the $j = 0$ waveform of Fig. 3. Solid (black), dashed (red), and dotted (blue) lines refer to the MP, KT, and LM algorithms, respectively.

VI. APPLICATION TO MERGER WAVEFORMS

Here we turn to the problem that motivates the present analysis of fitting algorithms for complex exponentials in noise. If the ringdown radiation emitted as a result of a binary merger shows no signs of nonlinearities or mode coupling, our fits should have a simple behavior: as we increase the starting time for the fit t_0, the black hole's

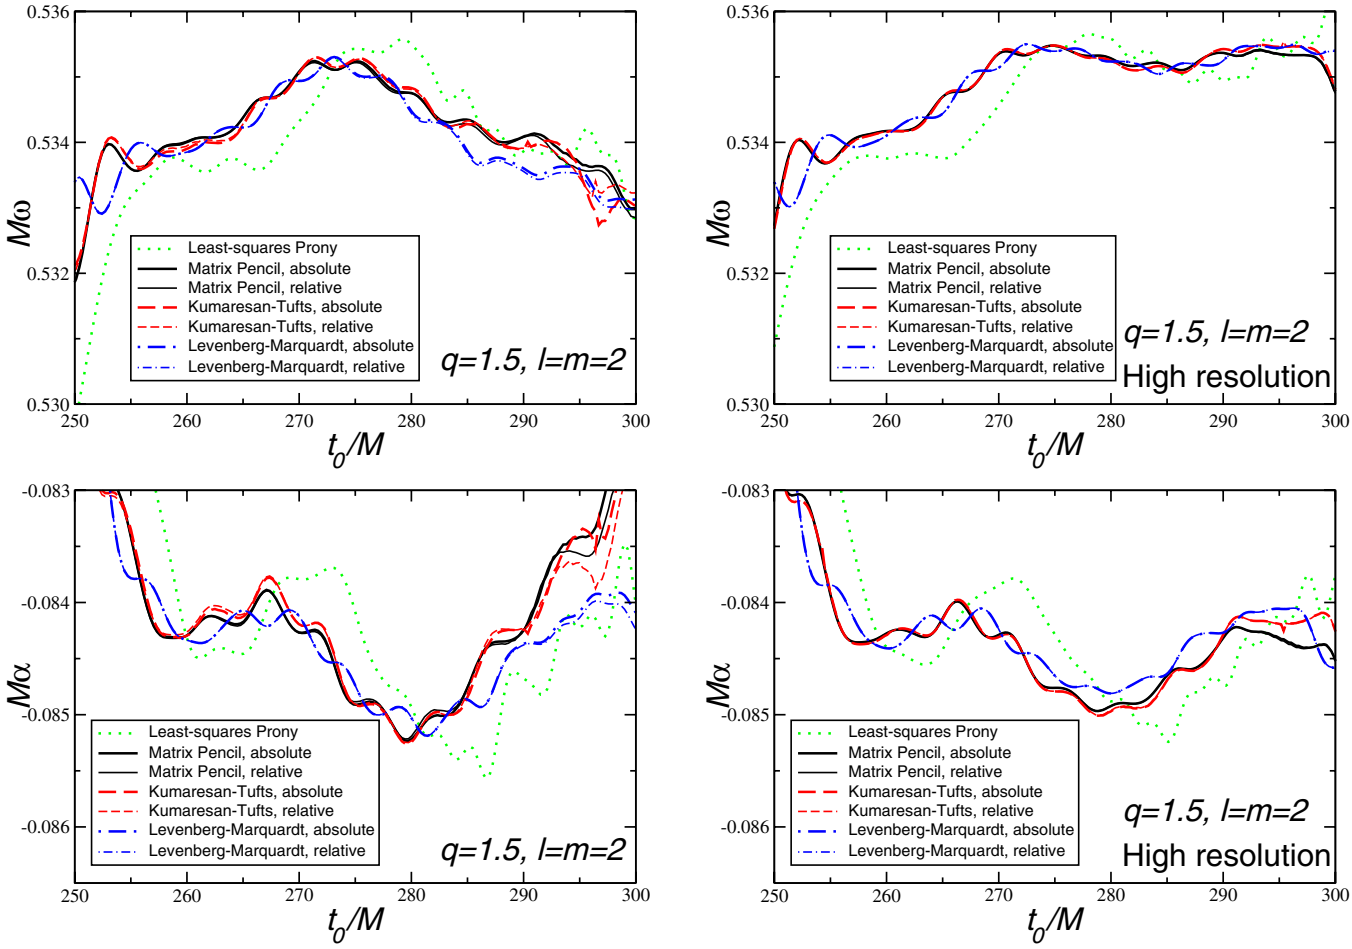

FIG. 5 (color online). Performance of different Prony methods in estimating the oscillation frequency ω (top) and damping factor α (bottom) for low and high-resolution runs (left and right, respectively). For concreteness we choose a binary with $q = 1.5$ and consider the fundamental mode with $l = m = 2$. Prony-type results refer to the complex waveform, while LM results are obtained by fitting the *real* part of the waveform only.

oscillation frequencies ω_{lmn} and damping factors α_{lmn} should *asymptotically approach a constant*, whose value is predicted by linear perturbation theory [17].

In Fig. 5 we show results of QNM fits performed on the $l = m = 2$ component of low-resolution and high-resolution merger simulations with mass ratio $q = 1.5$. Solid lines refer to the absolute truncation criterion of Sec. II, and dashed lines to the relative truncation criterion. Different truncation criteria affect the estimated parameters only for low-resolution simulations, and at late starting times ($t_0/M \gtrsim 290$). Results obtained fitting by a single complex exponential with standard least-squares Prony[8] are slightly off from the "best" fitting methods (MP, KT, and LM). Not surprisingly, there is remarkable agreement between KT and MP methods, and very good agreement between these two methods and the nonlinear least-squares LM fit.

The main conclusion to be drawn from Fig. 5 is that *all fitting routines consistently predict that frequencies and damping times have small (but nontrivial in structure) time variations*. Relative variations in ω are of order $\sim$0.5%, with the frequency reaching a local maximum at $t_0/M \sim 272$. Relative variations in α are slightly larger (roughly $\sim$2.5%), with $|\alpha|$ attaining a maximum (and the damping time correspondingly attaining a minimum) for $t_0/M \sim 280$. Notice also that increasing the resolution sensibly reduces the irregularities in the predicted frequency at all times, and produces a flattening of the estimated parameters for $280 \lesssim t_0/M \lesssim 300$. The decrease in ω and in $|\alpha|$ for $280 \lesssim t_0/M \lesssim 300$, that are visible in the left panels, are clearly an artifact of insufficient resolution. It would be interesting to perform Richardson extrapolation of the numerical results to determine if oscillations in the estimated parameters (which could be a sign of "new"

[8]Figs. 5 and 6 only display results for least-squares Prony obtained using the relative truncation criterion. However, we did check that using the absolute truncation criterion has a very small effect on the estimated parameters. We also checked that increasing the number of exponentials we search for does *not* improve the agreement of standard least-squares Prony with other methods.

physics) disappear in the limit of infinite resolution. We plan to address this problem in the near future. An analysis of the fine structure of the signal, and of its implications for gravitational wave phenomenology, will be presented in a forthcoming publication [30].

For the time being, we simply point out that such systematic variations in the oscillation frequency and/or damping time could be signs of nonlinearities and/or mode coupling in the numerical simulations. Variations could be due to the black hole's mass and angular momentum changing on a time scale which is longer than the QNM oscillation period, or to beating phenomena with other QNM frequencies. In our opinion, the fact that all fitting methods consistently predict the same "global" structure is convincing evidence that the existence of maxima and minima is *not* due to numerical errors in the fit.

In Fig. 6 we plot the quality factor of the oscillations as a function of t_0/M. The time variation of Q is basically dominated by the time variation of α. This is quite obvious, since $Q = \omega/(2\alpha)$ and time variations in α are larger than time variations in ω. The interest of plotting $Q(t_0/M)$ comes from the fact that, in linear theory and for $l = m$, Q is a monotonic function of j [17]. Roughly speaking, large variations in Q for the fundamental oscillation mode could mean that the Kerr angular momentum parameter is changing in time, or that there are significant corrections to the linear approximation.

Finally, in Fig. 7 we compare the performance of our two best fitting methods (MP and LM) in estimating oscillation frequencies and quality factors for different mass ratios: $q = 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0$. Deviations in Q between the two methods are quite significant, especially for large mass ratio, when the numerical simulations are less reliable. For $q = 4.0$ and low resolution, the LM method shows larger variations in ω than the MP method. This could be an effect of the larger bias in frequency of the LM method, as compared with the MP method (cf. Fig. 4). Once again, increasing the resolution produces a flattening of all curves, the effect being more pronounced for large mass ratios. A more systematic analysis of simulations for different mass ratios will appear elsewhere [30].

VII. CONCLUSIONS AND OUTLOOK

Our understanding of binary black hole mergers has recently been revolutionized by the success of numerical relativity simulations. A partial disappointment came from the remarkable simplicity of the inspiral-merger-ringdown transition: the nonlinearities of Einstein's theory do not seem to leave spectacular imprints on the merger, that seems to be a very short phase smoothly connecting the familiar "inspiral" and ringdown waveforms [12].

In this paper we argued that nonlinearities could show up in the *fine structure* of the ringdown waveform, as systematic time variations of the ringdown frequencies and damping times. We also showed by explicit calculations that such time variations *are* actually present in numerical waveforms (see Figs. 5–7). The variations we are looking for are typically ~100 times smaller than the QNM frequencies themselves. This smallness calls for accurate parameter estimation methods to extract ringdown frequencies from numerical simulations. We considered a class of well-studied and robust parameter estimation methods for complex exponentials in noise, which are modern variations of a linear parameter estimation technique first introduced in 1795 by Prony.

The comparison of different fitting methods can help resolve actual physical effects from systematic parameter estimation errors, due to the variance and bias of each particular fitting algorithm. For this reason we compared two variants of the original Prony algorithm (the Kumaresan-Tufts and matrix pencil methods [28,29]) against standard nonlinear least-squares techniques, such as the Levenberg-Marquardt algorithm. We found that the two classes of methods have comparable variance, but Prony-type methods tend to have slightly smaller bias.

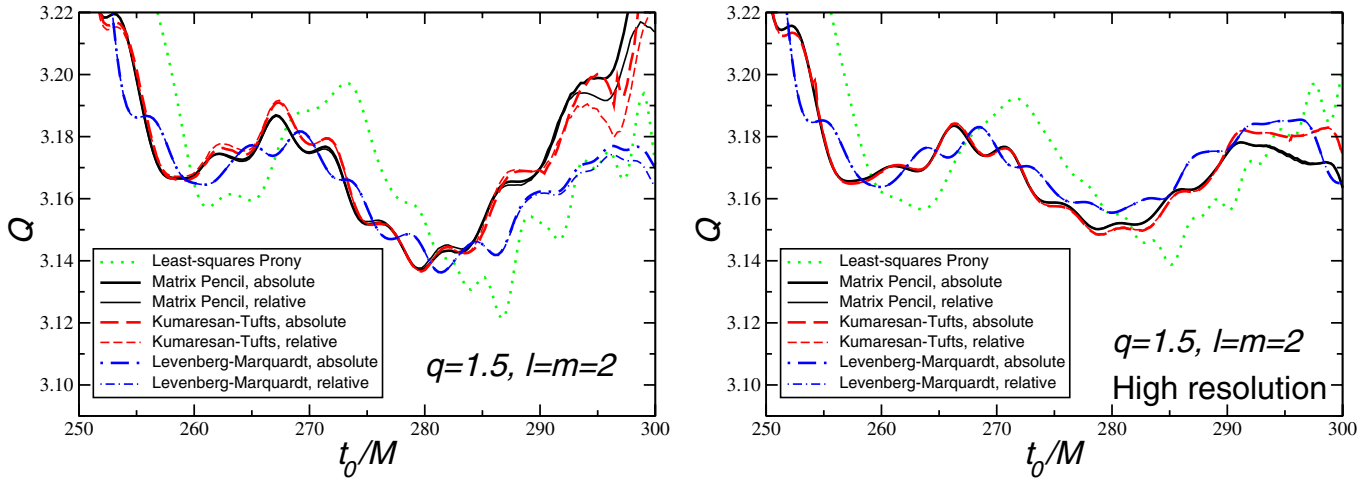

FIG. 6 (color online). Same as Fig. 5, but for the quality factor Q.

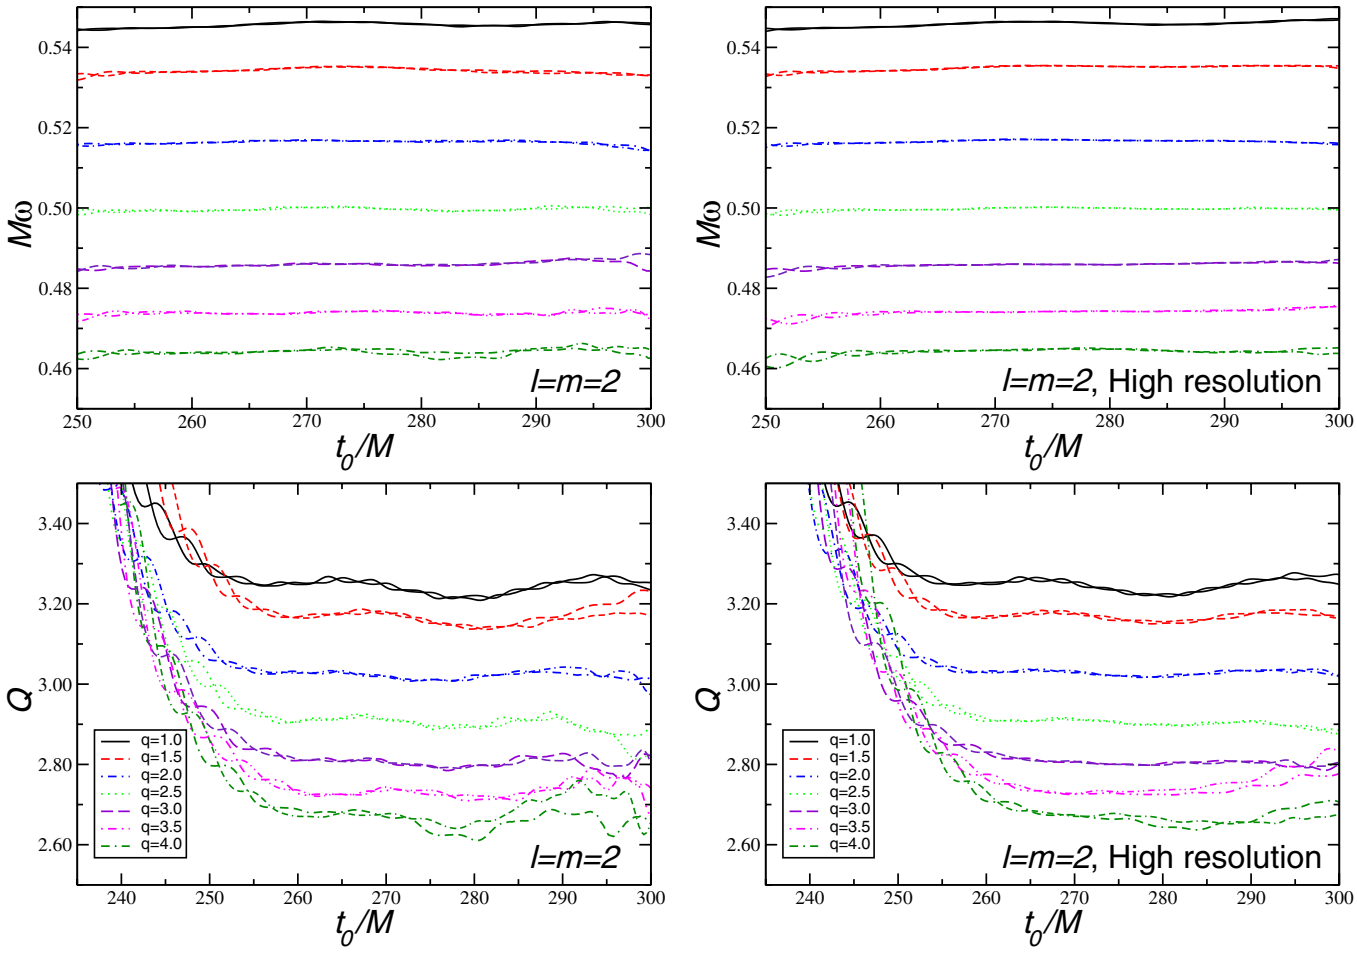

FIG. 7 (color online). Oscillation frequency (top) and quality factor (bottom) for modes with $l = m = 2$ and for low- and high-resolution runs (left and right, respectively), as a function of the starting time of the fit. Thick lines use MP, thin lines use LM. Values of the quality factor for $t/M \lesssim 250$ are often unphysically large. We truncate the waveform at late times when the absolute value of the amplitude drops below 10^{-4} (for low resolution) or 10^{-5} (for high resolution). In each panel, lines from top to bottom refer to $q = 1.0$, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0.

Prony methods have a number of advantages with respect to standard nonlinear least-squares techniques:

(1) They do not require an *initial guess* of the fitting parameters.

(2) They provide us with a simple, efficient way to estimate QNM frequencies for the *overtones*, and even to estimate how many overtones are present in the signal [47,48].

(3) They are explicitly designed to deal with *complex signals*, so they should be most useful for "generic" waveforms, such as those produced by spinning, precessing black hole binaries. In the case of non-trivial polarization of the gravitational radiation it might become crucial to fit simultaneously the real and imaginary parts of the signal to exploit optimally the information carried by both degrees of freedom.

(4) *Statistical properties* of Prony-based methods in the presence of noise (such as their variance and bias) are well studied and under control. Our Monte Carlo simulations suggest that the statistical properties of Prony-like methods as we vary the SNR are more consistent than for nonlinear least-squares methods, even when we consider a single damped sinusoid (the dotted blue lines in Fig. 4 have a very irregular behavior).

The methods introduced in this paper should be useful both theoretically and experimentally. From a theoretical standpoint, besides helping in the search for nonlinearities, Prony methods can also be used as "diagnostic tools" for numerical simulations. For example, monitoring the time variation of ringdown parameters from different multipolar components of the radiation we can check that the end-product of a merger *is* indeed consistent with the Kerr solution. An analysis of presently available numerical waveforms based on Prony methods is ongoing [30].

So far, Prony algorithms have generally been overlooked by the GW data analysis community. We believe that the techniques presented in this paper should prove useful to extract science from noisy ringdown signals in the (hope-

fully not too distant) future, when GW detection will finally become a reality.

ACKNOWLEDGMENTS

We are grateful to E. H. Djermoune, R. Kumaresan, S. Lawrence Marple, B. Porat, and G. Smyth for very useful correspondence and for sharing with us their numerical routines. We also thank B. Brügmann, J. Cardoso, M. Hannam, and S. Husa for useful discussions. V. C. acknowledges financial support from Fundação Calouste Gulbenkian through the Programa Gulbenkian de Estímulo à Investigação Científica. Computations were performed at HLRS (Stuttgart) and LRZ (Munich). This work was supported in part by DFG grant SFB/Transregio 7 ''Gravitational Wave Astronomy,'' by Fundação para a Ciência e Tecnologia (FCT)–Portugal through Project No. PTDC/FIS/64175/2006, by the National Science Foundation under Grant No. PHY 03-53180, and by NASA under Grant No. NNG06GI60 to Washington University. J. G. and U. S. acknowledge support from the ILIAS Sixth Framework Programme.

[1] B. Brügmann, W. Tichy, and N. Jansen, Phys. Rev. Lett. **92**, 211101 (2004).
[2] F. Pretorius, Phys. Rev. Lett. **95**, 121101 (2005).
[3] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, Phys. Rev. Lett. **96**, 111101 (2006).
[4] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter, Phys. Rev. Lett. **96**, 111102 (2006).
[5] F. Herrmann, D. Shoemaker, and P. Laguna, arXiv:gr-qc/0601026.
[6] U. Sperhake, arXiv:gr-qc/0606079.
[7] M. A. Scheel, H. P. Pfeiffer, L. Lindblom, L. E. Kidder, O. Rinne, and S. A. Teukolsky, Phys. Rev. D **74**, 104006 (2006).
[8] B. Brügmann, J. A. Gonzalez, M. Hannam, S. Husa, U. Sperhake, and W. Tichy, arXiv:gr-qc/0610128.
[9] J. A. Gonzalez, U. Sperhake, B. Brügmann, M. Hannam, and S. Husa, Phys. Rev. Lett. **98**, 091101 (2007).
[10] B. Szilagyi, D. Pollney, L. Rezzolla, J. Thornburg, and J. Winicour, arXiv:gr-qc/0612150.
[11] E. Berti, S. Iyer, and C. M. Will, Phys. Rev. D **74**, 061503 (2006).
[12] A. Buonanno, G. B. Cook, and F. Pretorius, arXiv:gr-qc/0610122 [Phys. Rev. D (to be published)].
[13] J. G. Baker, J. R. van Meter, S. T. McWilliams, J. Centrella, and B. J. Kelly, arXiv:gr-qc/0612024.
[14] K. D. Kokkotas and B. F. Schutz, Gen. Relativ. Gravit. **18**, 913 (1986); Mon. Not. R. Astron. Soc. **225**, 119 (1992).
[15] C. V. Vishveshwara, Nature (London) **227**, 936 (1970); S. Chandrasekhar and S. Detweiler, Proc. R. Soc. A **344**, 441 (1975); Kostas D. Kokkotas and Bernd G. Schmidt, Living Rev. Relativity **2**, 2 (1999), http://www.livingreviews.org/lrr-1999-2; H.-P. Nollert, Classical Quantum Gravity **16**, R159 (1999).
[16] E. W. Leaver, Proc. R. Soc. A **402**, 285 (1985).
[17] E. Berti, V. Cardoso, and C. M. Will, Phys. Rev. D **73**, 064030 (2006).
[18] Y. Zlochower, R. Gomez, S. Husa, L. Lehner, and J. Winicour, Phys. Rev. D **68**, 084014 (2003).
[19] E. Berti, V. Cardoso, and M. Casals, Phys. Rev. D **73**, 024013 (2006); **73**, 109902(E) (2006).
[20] E. N. Dorband, E. Berti, P. Diener, E. Schnetter, and M. Tiglio, Phys. Rev. D **74**, 084028 (2006).
[21] R. Boyer and J. Rosier, Proceedings of the 5th International Conference on Digital Audio Effects 2002; K. Hermusa, W. Verhelstb, P. Lemmerlingc, P. Wambacq, and S. Van Huffel, Signal Processing **85**, 163 (2005).
[22] M. Wax and T. Kailath, IEEE Trans. Acoust. Speech Signal Process. **33**, 387 (1985); J. J. Fuchs, IEEE Trans. Acoust. Speech Signal Process. **36**, 1846 (1988).
[23] A. I. Kulmentev, Eur. Phys. J. Appl. Phys. **25**, 191 (2004).
[24] E.-H. Djermoune, Ph.D. thesis, Centre de Recherche en Automatique de Nancy, Université Henri Poincarè, 2003.
[25] S. L. Marple, *Digital Spectral Analysis with Applications* (Prentice-Hall, New Jersey, 1987).
[26] R. N. McDonough and W. H Huggins, IEEE Trans. Autom. Control **13**, 408 (1968); A. G. Evans and R. Fischl, IEEE Transactions on Aerospace and Electronic Systems **aes-12**, 583 (1976).
[27] G. de Prony, J. l'École Polytechnique **1**, 24 (1795).
[28] R. Kumaresan and D. W. Tufts, IEEE Trans. Acoust. Speech Signal Process. **30**, 833 (1982); D. W. Tufts and R. Kumaresan, Proc. IEEE **70**, 975 (1982).
[29] Y. Hua and T. K. Sarkar, IEEE Trans. Acoust. Speech Signal Process. **38**, 814 (1990); T. K. Sarkar and O. Pereira, IEEE Trans. Antennas Propag. **37**, 48 (1995).
[30] E. Berti, V. Cardoso, J. A. Gonzalez, U. Sperhake, M. Hannam, S. Husa, and B. Brügmann, arXiv:gr-qc/0703053.
[31] J. G. Baker, M. Campanelli, C. O. Lousto, and R. Takahashi, Phys. Rev. D **65**, 124012 (2002).
[32] E. Berti and V. Cardoso, Phys. Rev. D **74**, 104020 (2006).
[33] R. H. Price, Phys. Rev. D **5**, 2419 (1972).
[34] K. Levenberg, Q. Appl. Math. **2**, 164 (1944); D. Marquardt, SIAM J. Appl. Math. **11**, 431 (1963).
[35] J. J. Moré, D. C. Sorensen, K. E. Hillstrom, and B. S. Garbow, in *Sources and Development of Mathematical Software*, edited by W. J. Cowell (Prentice-Hall, Englewood Cliffs, NJ, 1984), pp. 88–111.
[36] http://www.wolfram.com/products/mathematica/index.html.
[37] S. Marcos, *Les Méthodes à la Haute Résolution: Traitment Dántennes et Analyse Spectrale* (Hermès, Paris, 1998).
[38] M. A. Rahman and K.-B. Yu, IEEE Trans. Acoust. Speech Signal Process. **35**, 1440 (1987).

[39] A. C. Kot, S. Parthasarathy, D. W. Tufts, and R. J. Vaccaro, in Proceedings of the International Conference on Acoustics, Speech and Signal Processing (1987), p. 1549.
[40] B. Porat and B. Friedlander, IEEE Trans. Acoust. Speech Signal Process. **35**, 231 (1987); Y. Hua and T. K. Sarkar, IEEE Trans. Acoust. Speech Signal Process. **36**, 228 (1988); A. Okhovat and J. R. Cruz, in Proceedings of the International Conference on Acoustics, Speech and Signal Processing (1989), p. 2286.
[41] E.-H. Djermoune and M. Tomczak, 12th European Signal Processing Conference (EUSIPCO'04), Wien, Austria, 2004.
[42] B. Porat, *Digital Processing of Random Signal: Theory and Methods* (Prentice-Hall, Englewood Cliffs, NJ, 1993).
[43] J. C. Hamann and S. Singh, http://www.engr.uconn.edu/sas03013/pronytool/prony.html.
[44] C. Pozrikidis, http://dehesa.freeshell.org/SOFT/FDLIB/FDLIB/08_stab/prony/.
[45] G. Smyth, http://www.statsci.org/other/prony.html.
[46] J. H. McClellan, C. Sidney Burrus, A. V. Oppenheim, T. W. Parks, R. W. Schafer, and H. W. Schuessessler, *Computer Based Exercises for Signal Processing using MATLAB* (Prentice-Hall, Englewood Cliffs, NJ, 1997).
[47] V. U. Reddy and L. S. Biradar, IEEE Trans. Signal Process. **41**, 2872 (1993).
[48] M. Wax and T. Kailath, IEEE Trans. Acoust. Speech Signal Process. **33**, 387 (1985).
[49] E.-H. Djermoune and M. Tomczak, 11th European Signal and Image Processing Conference (EUSIPCO'02), Toulouse, France, 2002, p. 139.
[50] G. E. P. Box and M. E. Muller, Ann. Math. Stat. **29**, 610 (1958).
[51] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in Fortran, Second Edition* (Cambridge University Press, Cambridge, England, 1992).