

Longitudinal electron bunch profile reconstruction by performing phase retrieval on coherent transition radiation spectra

S. I. Bajlekov,¹ M. Heigoldt,² A. Popp,² J. Wenz,² K. Khrennikov,² S. Karsch,² and S. M. Hooker^{1,*}

¹*Department of Physics, University of Oxford, OX1 3PU, United Kingdom*

²*Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Straße 1, D-85748 Garching, Germany*

(Received 18 August 2012; published 23 April 2013)

The application of phase-retrieval algorithms to the reconstruction of the longitudinal bunch profile of an electron bunch from the spectrum of the coherent transition radiation (CTR) it produces is considered. The development of a new algorithm for this application, the Bubblewrap algorithm, is described. Tests with synthetic data show successful reconstruction of the longitudinal profile of single and double electron bunches, provided that the CTR spectrum is known over a sufficiently wide range. The application of the Bubblewrap algorithm to the reconstruction of laser-accelerated electron bunch profiles from experimental data is demonstrated. The results are shown to be consistent with estimates of the bunch length obtained by other methods.

DOI: [10.1103/PhysRevSTAB.16.040701](https://doi.org/10.1103/PhysRevSTAB.16.040701)

PACS numbers: 41.60.Dk, 42.30.Rx, 52.38.Kd, 41.75.Ht

I. INTRODUCTION

When a charged particle crosses a boundary between media of different dielectric indices it causes emission of very broadband transition radiation [1]. This transition radiation can be used to gain information about the longitudinal and transverse profiles of particle beams, with essentially no disruption to relativistic beam propagation. Nondestructive diagnostics of particle beams—particularly those which operate in a single-shot regime—are invaluable for understanding the operation of the accelerator and for providing feedback to the accelerator control system. An area of significant current interest is the development of single-shot diagnostics for the femtosecond-duration electron bunches generated by laser-wakefield accelerators. The physics of laser-wakefield acceleration ensures that the accelerated bunches have a duration equal to a fraction of the plasma period, i.e., of order 10 fs or less [2–4], which is too short for established techniques, such as transverse-deflecting structures [5,6]. Meanwhile, the potential applications of these bunches, for instance as drivers of compact free-electron lasers [7], depend critically on the bunch duration. It is therefore important to develop new diagnostics for these ultrashort bunches.

The emission of incoherent transition radiation at a thin screen placed in the path of a particle bunch can be used to determine the transverse profile of the bunch [8,9]. Additionally, transition radiation emitted at wavelengths longer than the bunch length exhibits a degree of coherence, due to the fact that emission from different particles

occurs approximately in phase—this is so-called coherent transition radiation (CTR). The onset of this coherence is reflected in the transition radiation spectrum, with (coherent) long-wavelength radiation being more intense than (incoherent) short-wavelength radiation by a factor equal to the number of particles emitting coherently. Ignoring transverse effects, the spectral intensity of the emitted radiation is proportional to the Fourier transform of the longitudinal bunch profile,

$$\frac{d^2U}{d\omega d\Omega} = \frac{d^2U_{GF}}{d\omega d\Omega} N_e |\tilde{\rho}_{\parallel}(\omega)|^2, \quad (1)$$

where U is the emitted energy, ω is the angular frequency, Ω is the solid angle subtended at the observation screen, N_e is the total number of electrons, and $\tilde{\rho}_{\parallel}(\omega) = \mathcal{F}[\rho_{\parallel}(t)]$ is the Fourier transform of the longitudinal bunch profile $\rho_{\parallel}(t)$, and where

$$\frac{d^2U_{GF}}{d\omega d\Omega} = \frac{e^2}{4\pi^3 \epsilon_0 c} \frac{\beta_{\gamma}^2 \sin^2 \theta}{(1 - \beta_{\gamma}^2 \cos^2 \theta)^2} \quad (2)$$

is the Ginzburg-Frank formula for the spectral energy emitted by a single particle with normalized velocity $\beta_{\gamma} = v/c$ as a function of emission angle θ .

The spectrum of the radiation therefore contains information on the longitudinal bunch profile. Measurements of the autocorrelation of the CTR by scanning interferometers have been used to estimate the bunch length at several conventional accelerator facilities [10,11], and recently Wesch *et al.* have described a single-shot, broadband multigrating spectrometer for measurement of the CTR spectrum in the ranges 5–44 μm and 45–435 μm [12]. Estimates of the duration of the electron bunches produced by laser-driven plasma accelerators have been obtained on the basis of the CTR spectrum by, among others, Leemans *et al.* [2], Faure *et al.* [3], and Glinec *et al.* [4].

*simon.hooker@physics.ox.ac.uk.

Published by the American Physical Society under the terms of the [Creative Commons Attribution 3.0 License](https://creativecommons.org/licenses/by/3.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

More recently, Lundh *et al.* [13] carried out a measurement of the CTR at multiple wavelengths that allowed the establishment of rough constraints on the longitudinal bunch profile.

Direct reconstruction of the longitudinal profile $\rho_{\parallel}(t)$ from an inverse Fourier transform of the measured CTR spectrum is not possible since: (i) a CTR measurement only yields the amplitude of $\tilde{\rho}_{\parallel}(\omega)$, and not its phase; and (ii) in practice the CTR spectrum can only be measured over a limited range of ω . In these circumstances an estimate of the duration of the bunch can be obtained by assuming a parametrized shape for the longitudinal profile and fitting the CTR spectrum calculated from this to the measured spectrum.

Alternatively, the longitudinal profile may be deduced from a Kramers-Kronig analysis and by making the minimum phase assumption, as first suggested by Lai and Sievers [14,15] and subsequently used in the analysis of many measurements [14–18]. The Kramers-Kronig method suffers from several limitations. The method requires the spectrum to be known for all frequencies, and hence in practice it is necessary to extrapolate the spectrum through spectral regions in which measurements are not made, especially the region near zero frequency. Further, without additional knowledge, or assumptions—such as a parametrized longitudinal profile—only the so-called minimal spectral phase can be retrieved [14,15]. However this may be insufficient, as in some (realistic) circumstances additional phase contributions, known as the Blaschke phase, can change the deduced longitudinal profile significantly. Given these limitations it is interesting to explore alternative methods for reconstructing the longitudinal profile from full or partial measurements of the radiation spectrum.

Provided certain conditions are met, phase-retrieval algorithms can find the inverse transform of $|\tilde{\rho}_{\parallel}(\omega)|$ even in the absence of phase information, and when the amplitude is not known for all ω [19]. Algorithms of this type have been employed for this purpose across many scientific fields, such as reconstruction of crystal structure from x-ray diffraction patterns [20], and determination of the modes excited in a waveguide from measurements of the diffraction of the emerging beam [21].

In this paper we consider, for the first time, the application of phase-retrieval algorithms to deducing the longitudinal bunch profile of a charged particle beam from measurements of part of the CTR spectrum. The paper is structured as follows. In Sec. II we provide a brief overview of phase-retrieval algorithms developed to date. In Sec. III we present a new algorithm—which we call the Bubblewrap algorithm—that is specifically targeted at longitudinal bunch profile reconstruction. In Sec. IV, we use synthetic data to demonstrate the ability of the Bubblewrap algorithm to reconstruct the bunch profiles for several cases, and also outline its limitations. We then transition

to the real-world application of the Bubblewrap algorithm. In Sec. V we describe a recent experiment where broadband single-shot measurements of CTR spectra emitted by electron bunches from a laser-wakefield accelerator were conducted. Finally, in Sec. VI we present the reconstruction of longitudinal electron bunch profiles through application of the Bubblewrap algorithm to CTR spectra obtained in this experiment, and compare the results with fits to an assumed bunch profile.

II. PHASE-RETRIEVAL ALGORITHMS

Given the square modulus of the Fourier transform of a signal, the missing phase information (and the signal itself) can under certain circumstances be retrieved *a posteriori*. This concept is underpinned by a note by Sayre [22], who pointed out that a function with finite support is uniquely defined by knowledge of its Fourier transform at a limited number of points. In particular, a function with support in real space over the range $x \in [-a/2, a/2]$ is completely specified by the values of its Fourier transform at the points $k = 0, \pm 2\pi/a, \pm 4\pi/a, \dots$. Hence, if the amplitude of the transform is oversampled it is possible to use the additional data to deduce the phase of the limited set of points required for reconstruction of the function; in many experiments the transform is indeed oversampled.

The first practical phase-retrieval algorithm was introduced in 1972 by Gerchberg and Saxton [23], for the purpose of reconstructing the 2D profile of an object from its diffraction pattern. The Gerchberg-Saxton algorithm (hereafter the GS algorithm) is iterative, as are all known phase-retrieval algorithms.

A. Gerchberg-Saxton algorithm

We first present the Gerchberg-Saxton algorithm, as its operation informs that of others discussed later. Our treatment follows that of Fienup [19]. We introduce the function $f(x)$ that we would like to reconstruct from measurements of the modulus of its Fourier transform, $|F(k)| = |\mathcal{F}[f(x)]|$. The function $g_n(x)$ is our estimate of $f(x)$ at the n th iteration, and $G_n(k)$ is its Fourier transform. Figure 1 illustrates the four-step iteration cycle.

The algorithm continually switches between Fourier space and real space, and ensures that the candidate estimate g matches the constraints present in each domain. In the Fourier domain the only constraint is that the modulus of G must equal the measured amplitude of the spectrum, $|F|$. Since there is no phase information about F , the phase of $G_n(k)$, $\psi_n(k)$, is retained. The modulus of G is set equal to $|F|$ only for the points for which measurements exist; for other points the iterative estimate G is preserved. According to the formalism utilized by Elser [24], this— together with the forward and inverse Fourier transforms— constitutes a “projection,” π_{mod} of g onto the closest function that satisfies the Fourier-space modulus constraint, $g' = \pi_{\text{mod}}[g]$.

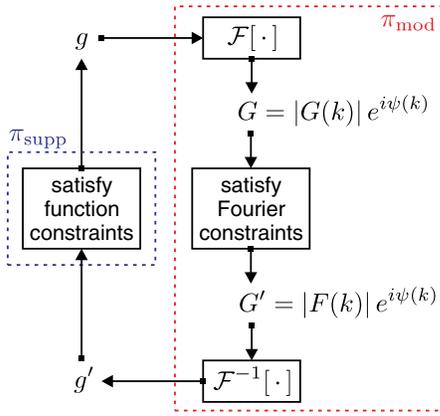


FIG. 1. Outline of the Gerchberg-Saxton phase-retrieval algorithm. For further details on each step see the text. Dashed boxes indicate the steps that form projections of a candidate function onto the support constraints, π_{supp} , and the Fourier-space modulus constraints, π_{mod} , respectively.

The real-space constraints can be more diverse according to the problem to be solved. It is common to require $f(x)$ to be zero outside a certain support. Additionally, as in the present application, $f(x)$ may be required to be real and positive inside the support. We note that the requirement for $f(x)$ to be real is automatically satisfied for all iterates g_n if the initial estimate g_0 is real valued. A set of points γ can therefore be defined for which $g'_k(x)$ violates the real-space constraints, and for which the function's value is set to zero:

$$g_{k+1}(x) = \begin{cases} g'_k(x), & x \notin \gamma, \\ 0, & x \in \gamma. \end{cases} \quad (3)$$

In Elser's formalism, this constitutes the support projection, π_{supp} .

The modifications to g and G described above are the *minimum* changes to the functions that are required to make them comply with the constraints in each domain. As discussed by Fienup, each iteration reduces the error in the estimate, where the error at the k th iteration is defined as

$$E_k = \sqrt{\sum_x |g_{k+1}(x) - g'_k(x)|^2}. \quad (4)$$

The GS algorithm is therefore also known as the error-reduction algorithm. However, although the error decreases at each iteration, the algorithm is prone to stagnation. After a number of iterations convergence can become very slow, and sometimes the candidate function can converge towards a local minimum [19]. This is in part a symptom of the fact that the algorithm has a tendency to "explore" a relatively small set of candidate functions with respect to the starting point $g_0(x)$.

B. Hybrid input-output algorithm

As first described by Fienup [19], the Gerchberg-Saxton algorithm can be enhanced in several ways. One of them involves recognizing that $g_k(x)$ does not need to be the current best estimate of $f(x)$, but can instead be considered as the driving function for the next input, $g'_k(x)$. In this sense the left-hand side of the procedure in Fig. 1 need not force $g(x)$ to satisfy the actual real-space constraints, but can instead be an arbitrary manipulation, such that the next $g'(x)$ will be closer to convergence. This approach forms a class of algorithms known as "input-output" algorithms.

Fienup compares several input-output algorithms, as well as a number of others, against the Gerchberg-Saxton algorithm. Empirically, the best-performing appears to be the hybrid input-output (HIO) algorithm, in which the input function $g_{k+1}(x)$ is obtained from the output $g'_k(x)$ via

$$g_{k+1}(x) = \begin{cases} g'_k(x), & x \notin \gamma, \\ g_k(x) - \beta g'_k(x), & x \in \gamma, \end{cases} \quad (5)$$

where β is a free parameter. Fienup's numerical tests on reconstruction of an object from a 2D diffraction pattern indicate best performance for $\beta \approx 1$. The hybrid input-output algorithm has been applied successfully for reconstruction of object data from a range of diffraction patterns, notably by Marchesini *et al.* [25].

C. Difference map algorithms

Elser showed [24,26] that the Gerchberg-Saxton algorithm consists of alternating application of the projections π_{mod} and π_{supp} , i.e. $g_k = \pi_{\text{mod}} \circ \pi_{\text{supp}} \cdots \circ \pi_{\text{mod}} \circ \pi_{\text{supp}}[g_0]$. The same author has described another class of algorithms which utilize difference maps to facilitate convergence to a fixed point, i.e., a function f that remains unchanged under either of the two projections. The optimum difference map is found to be

$$g_{k+1} = g_k + \beta_{\text{dm}} \{ \pi_{\text{supp}} [(1 + \beta_{\text{dm}}^{-1}) \pi_{\text{mod}}[g_k] - \beta_{\text{dm}}^{-1} g_k] - \pi_{\text{mod}} [(1 - \beta_{\text{dm}}^{-1}) \pi_{\text{supp}}[g_k] + \beta_{\text{dm}}^{-1} g_k] \}, \quad (6)$$

where β_{dm} is a free parameter. Elser selects $\beta_{\text{dm}} = 1$. We note that the projection performed by the difference map algorithm with $\beta_{\text{dm}} = 1$ is equivalent to the hybrid input-output algorithm with $\beta = 1$.

D. Choice of function support and initial estimate

The initial estimate, $g_0(x)$, is generally taken to be a randomly generated real function—not necessarily compliant with the support constraint—in order not to bias the phase-retrieval algorithm towards a particular output. This approach allows an algorithm to be run multiple times for different initial g_0 in order to obtain a set of outputs that match all constraints, and to correspondingly verify the uniqueness—or otherwise—of the reconstructed function.

On the other hand, an appropriate choice of support is crucial to the successful execution of a phase-retrieval algorithm. It has to be large enough to accommodate the entire reconstructed function, while being sufficiently restrictive for quick convergence to an unambiguous target function to occur. In particular, an excessively large support may yield significantly differing results for runs with different starting points.

The support constraints can be chosen on the basis of additional information: for example, in the case of x-ray diffraction, a low-resolution image of the object obtained using conventional optics or scanning electron tomography can inform the choice of support. They can also be chosen on the basis of physical considerations: e.g., the length of an electron bunch from a laser-wakefield accelerator is expected to be limited by the bubble size, or plasma wavelength. In addition, the autocorrelation function of the object can be used as an indication of the outer limits of its extent for the purposes of reconstruction. Since, by the Wiener-Khinchin theorem, the autocorrelation $f(x) \star f(x) = \mathcal{F}^{-1}[|F|^2]$, this estimate of the support can be calculated from the measured spectrum $|F|^2$, provided the spectrum is known over a sufficiently broad range of frequencies. Marchesini *et al.* [25] give further details on possible choices of support.

E. Adaptive support selection: The Shrinkwrap algorithm

In light of the importance of the choice of function support, a key recent development has been the suggestion by Marchesini *et al.* to adaptively change the support during the reconstruction process, shrinking it around subsequent iterates of the candidate function [20]. In this approach the initial support is obtained by applying a threshold to the autocorrelation of the diffraction pattern (or spectrum). The HIO algorithm is then applied for 20 iterations with $\beta = 0.9$. At that point the candidate function is convolved with a Gaussian of rms σ , and a threshold is applied at 20% of its maximum in order to obtain a new support. After another 20 iterations of the HIO algorithm the support is recalculated again by the same means, and so forth. The Gaussian used for blurring has an initial width of $\sigma = 3$ pixels, and is decreased by 1% every 20 iterations until it reaches 1.5 pixels. As a consequence, each time the support is recalculated it follows the outline of the candidate function more and more closely. For this reason the algorithm is known as the ‘‘Shrinkwrap’’ algorithm.

The Shrinkwrap algorithm has been applied successfully to a range of diffraction-related problems, demonstrating the possibility for successful reconstruction without *a priori* information [20,27]. However, we found its present form inadequate for the reconstruction of bunch profiles from hypothetical or real CTR spectra. We believe this to be due to the fundamentally different number of dimensions: two in the case of diffraction data to which

Shrinkwrap was applied, and one in the case of transition radiation data. In the following section we describe a new algorithm that combines several of the concepts reviewed here and is tailored to reconstruction of longitudinal bunch shapes from CTR spectra.

III. AN ALGORITHM FOR RECONSTRUCTION OF BUNCH PROFILES FROM CTR DATA

In order to select the most appropriate algorithm for reconstructing longitudinal bunch profiles from CTR spectra, the GS, HIO, and difference map algorithms were applied to simulated and measured CTR spectra. For the HIO and the difference map algorithms, a range of values of β and β_{dm} were trialed. Several approaches for adaptive support selection were also tested. Here we present the results of this survey, and outline the approach that appears to yield the best performance under a broad range of circumstances.

A. Selection of algorithm

Preliminary tests were carried out with synthetic longitudinal bunches with Gaussian profiles of rms length ranging from 0.5 to 3 μm . The corresponding simulated spectral data was truncated for wavelengths longer than 7 μm in order to emulate realistic reconstruction from experimental data (see Sec. V). The GS algorithm was found to perform well for a number of iterations, but it tended to stagnate or produce incorrect reconstructions—for example multiple bunches—whenever the support allowed for such ambiguity.

The HIO algorithm, with $\beta = 1$, was found to probe a broader range of candidate functions. However, it also tended to produce rather large changes in the candidate function from one iteration to the next, even when the function appeared to be close to a suitable solution. It was found that a compromise could be reached by reducing β , which correspondingly reduced the change in the candidate function at each iteration [cf. Eq. (5)]. Figure 2 illustrates the variation in the final error—as defined by Eq. (4)—and the rms length of the reconstructed function after 990 iterations of the HIO algorithm followed by ten iterations of the GS algorithm, performed on input data corresponding to the CTR spectrum generated by a single Gaussian of rms length 2 μm , on a support of length 15 μm .

Ten reconstruction runs were carried out for each value of β , which was varied between 0.1 and 1.2 in steps of 0.1. For these conditions, the HIO algorithm appeared to perform best for $\beta \approx 0.3$ –0.4. This difference in the optimum value of β compared with the previously reported [19,25] optimum value of $\beta \approx 1$ appears to be a reflection of the difference between 1D and 2D reconstructions. It also highlights the trade-off between rapid exploration of a large range of functions achieved at larger values of β ,

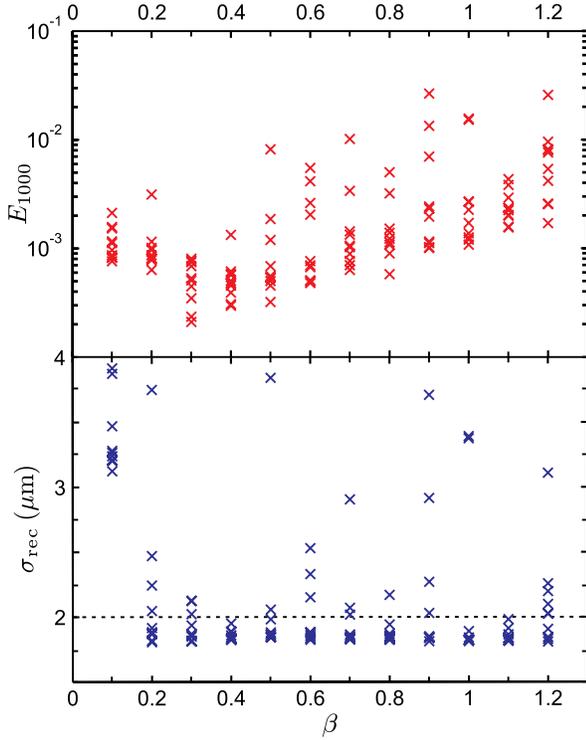


FIG. 2. Performance of the hybrid input-output algorithm for different values of the parameter β , when reconstructing the spectrum of a Gaussian of $2 \mu\text{m}$ rms length on a fixed support of length $15 \mu\text{m}$. The rms length of the reconstructed profiles is σ_{rec} , while E_{1000} shows the reconstruction error after 1000 iterations, as defined by Eq. (4). The spectrum is truncated at wavelengths above $7 \mu\text{m}$, leading to consistent underestimates of σ_z . Each reconstruction run consists of 990 iterations of the HIO algorithm followed by ten iterations of the Gerchberg-Saxton algorithm. Ten reconstructions were carried out for each value of β .

and the ability to accurately converge on the target function at lower values.

As reported by Fienup [19], we found that performing a small number of iterations of the GS algorithm after a sequence of HIO iterations helps minimize the error. Application of the GS algorithm was particularly good at eliminating nonzero values of the candidate function outside the support—i.e., forcing it to satisfy the support constraint—without significantly affecting the values inside the support.

For the synthetic bunch profile chosen here, the reconstructed function generally has an rms length somewhat less than the original $2 \mu\text{m}$ due to the fact that we allowed all spectral components with $\lambda > 7 \mu\text{m}$ to vary freely, as would be the case when experimental data is unavailable beyond $7 \mu\text{m}$. Under ideal conditions, a Gaussian bunch of rms length σ_z emits CTR with an intensity spectrum that varies as $\exp(-[\omega\sigma_z/c]^2)$, where $\omega = 2\pi c/\lambda$ is the angular frequency. While high-frequency components of the spectrum contain information about short-range longitudinal structure in the bunch, in order to make an accurate

estimate of the overall bunch length it is necessary to have spectral information at frequencies *less* than approximately the rms width of the spectrum. Hence, accurate determination of the bunch length requires measurements of the spectrum for frequencies $\omega \lesssim \omega_c \equiv c/\sigma_z$ or $\lambda \gtrsim \lambda_c \equiv 2\pi\sigma_z$. The reconstructions of synthetic data described here and in Sec. IV confirm this requirement empirically.

Finally, we found that the difference map algorithm suffered from large iteration-to-iteration changes similar to those that impeded the HIO algorithm's ability to converge on a solution. This is natural, as the two algorithms are equivalent for $\beta_{\text{dm}} = \beta = 1$. For the HIO algorithm this issue could be addressed by reducing β , yet for the difference map algorithm it was found that varying the parameter β_{dm} altered the algorithm's behavior [24], but did not appear to improve its performance with respect to several synthetic scenarios involving CTR data.

On the basis of this survey we selected the HIO algorithm as the primary means of phase retrieval, noting the benefit of applying the Gerchberg-Saxton algorithm at the final stages of reconstruction.

B. Adaptive support selection

A suitable support is crucial for a successful reconstruction. We found that a Gaussian of rms length $1 \mu\text{m}$ is reconstructed correctly after 990 iterations of the HIO algorithm followed by ten iterations of the GS algorithm in virtually every case on a support of length $10 \mu\text{m}$, but almost never on a support of length $15 \mu\text{m}$, independent of the value of β . Obtaining an initial support from the spectrum's autocorrelation function is not possible in the present application since there is no knowledge of the long-wavelength components. In addition, there is no sufficiently accurate *a priori* knowledge of the approximate bunch length; deducing a support on the basis of physical considerations—such as setting the plasma wavelength as a limit on the bunch length—is unlikely to be sufficiently restrictive for robust reconstruction.

It therefore appears beneficial to calculate the support adaptively during the reconstruction, as in the Shrinkwrap algorithm. The approach to support selection used by the Shrinkwrap algorithm, as described in Sec. II E, tends to lead to reconstructed profiles that are consistently *shorter* than the synthetic profiles used to generate input data. This is a natural outcome of the algorithm's tendency to restrict the size of the support as much as possible. While this approach may be valid for reconstructing the images of collections of compact and possibly scattered objects, like those considered in Ref. [20], it does not appear to be suitable for accurate bunch profile reconstruction in one dimension. We therefore modify the approach, as described below.

C. The Bubblewrap algorithm

In the original Shrinkwrap algorithm the threshold level applied to the candidate function each time the support was

recalculated was fixed at 20% of the peak. In the case of CTR data, a threshold at this level results in a support that cannot accommodate any low-current tails in the electron bunch profile, although these are likely to exist in reality. We find that a suitable way to address this is by starting at a high threshold level, which is subsequently reduced at each recalculation of the support. This ensures that the function reconstructed during the initial stages of the phase-retrieval process is still as compact as possible, however during the latter stages more subtle details are allowed to emerge near the edges. Furthermore, we find it beneficial initially to run the HIO algorithm with a high value of β —which allows a wider range of candidate functions to be probed—but subsequently to reduce β as the function converges towards a self-consistent profile, thus preventing excessive changes of the function.

In recognition of its relation to the Shrinkwrap algorithm, we call this approach the Bubblewrap algorithm. After trialling different combinations of parameters, the following approach appears to yield good results for a range of inputs.

(i) The initial support size is set to 20 μm . This can shrink or grow during the process, so the initial value is not crucial.

(ii) Between each recalculation of the support, the HIO algorithm is run for 45 iterations, and the Gerchberg-Saxton algorithm is run for five iterations. We will call this set of 50 iterations an *iteration cycle*.

(iii) The reconstruction process is run for a total of 200 iteration cycles (10 000 iterations in total).

(iv) The parameter β for the HIO algorithm is initially set equal to 1.0, and decreased by 1.5% after each iteration cycle, reaching a final value of 0.05.

(v) The initial threshold for support recalculation is set to 25%, and decreased by 2.5% (in relative terms) after each iteration cycle, bringing it down to a final value of 0.16%. The threshold is applied with respect to the peak value of the reconstructed function at the end of each iteration cycle.

(vi) Before calculating the support the function is smoothed by a Gaussian function of rms length σ . The parameter σ is initially set equal to 3 data points, and is decreased by 0.5% after each iteration cycle so that after 200 iteration cycles it reaches a final value of 1.1 data points.

These parameters are not intended to be universally applicable: some phase retrievals may require more iterations, and some may reach the target function after much fewer. However, they form a suitable starting point, and the parameters can be adjusted in light of the algorithm's performance. The number of data points used in the reconstructions presented here is $2^{13} = 8192$, however a significantly smaller number can be used without substantively affecting the effectiveness of the phase-retrieval process. This is discussed further in Sec. VI.

IV. RECONSTRUCTION OF PROFILES FROM SYNTHETIC SPECTRAL DATA

The Bubblewrap algorithm was applied to several synthetic spectra calculated from various assumed bunch shapes. The synthetic spectra were truncated at wavelengths longer than 7 μm , corresponding to a frequency of $2.7 \times 10^{14} \text{ rad s}^{-1}$; the signal at frequencies below this value was allowed to vary freely during the reconstruction. Since phase retrieval cannot fix the absolute position or the direction of the reconstructed function, in order to compare the reconstructed profiles they were shifted so that their peaks overlap, if necessary inverted, and rescaled to the same peak value. Ten reconstructions were carried out for each input spectrum, in order to ascertain the stability of the reconstruction.

For single Gaussian bunch shapes, we found that the lack of spectral information at long wavelengths impedes the accurate reconstruction of longer bunch profiles, leading to underestimates of their length. However, even in these cases the reconstructed bunch shape was still approximately Gaussian, occasionally with slight asymmetries. Table I summarizes the accuracy of reconstruction for several sample cases. We see that for bunches of length $\geq 2 \mu\text{m}$ the deviation in reconstructed bunch length becomes significant. For the case of the 3 μm -rms profile, in some cases the phase retrieval converges on a profile containing a pair of peaks rather than a single one, as reflected in the figures shown in the table. These findings are in reasonable agreement with the prediction of Sec. III A that accurate reconstruction of a bunch of length σ_z requires spectral information at wavelengths above $\lambda_c = 2\pi\sigma_z$.

Figure 3 shows sample reconstructions of asymmetric Gaussian profiles. Again, it is seen that the reconstruction becomes inaccurate whenever important spectral components of the original profile are left as free parameters during the phase retrieval. Figures 4 and 5 show the reconstruction of double Gaussian bunches, in which the two Gaussian profiles are either overlapping or separate. The corresponding figure captions give more detail on the reconstructed profiles.

A subset of the above reconstructions was also performed after introducing Gaussian noise to the synthetic

TABLE I. Phase retrieval for synthetic Gaussian bunch data. The given error margins are standard deviation over the ten reconstructions (or subsets thereof, in the case of 3 μm).

Original σ_z (μm)	Retrieved σ_z (μm)
0.5	0.500 ± 0.001
1.0	0.978 ± 0.004
1.5	1.414 ± 0.006
2.0	1.736 ± 0.019
3.0	2.066 ± 0.053 (7/10 cases), 4.369 ± 0.270 (3/10 cases)

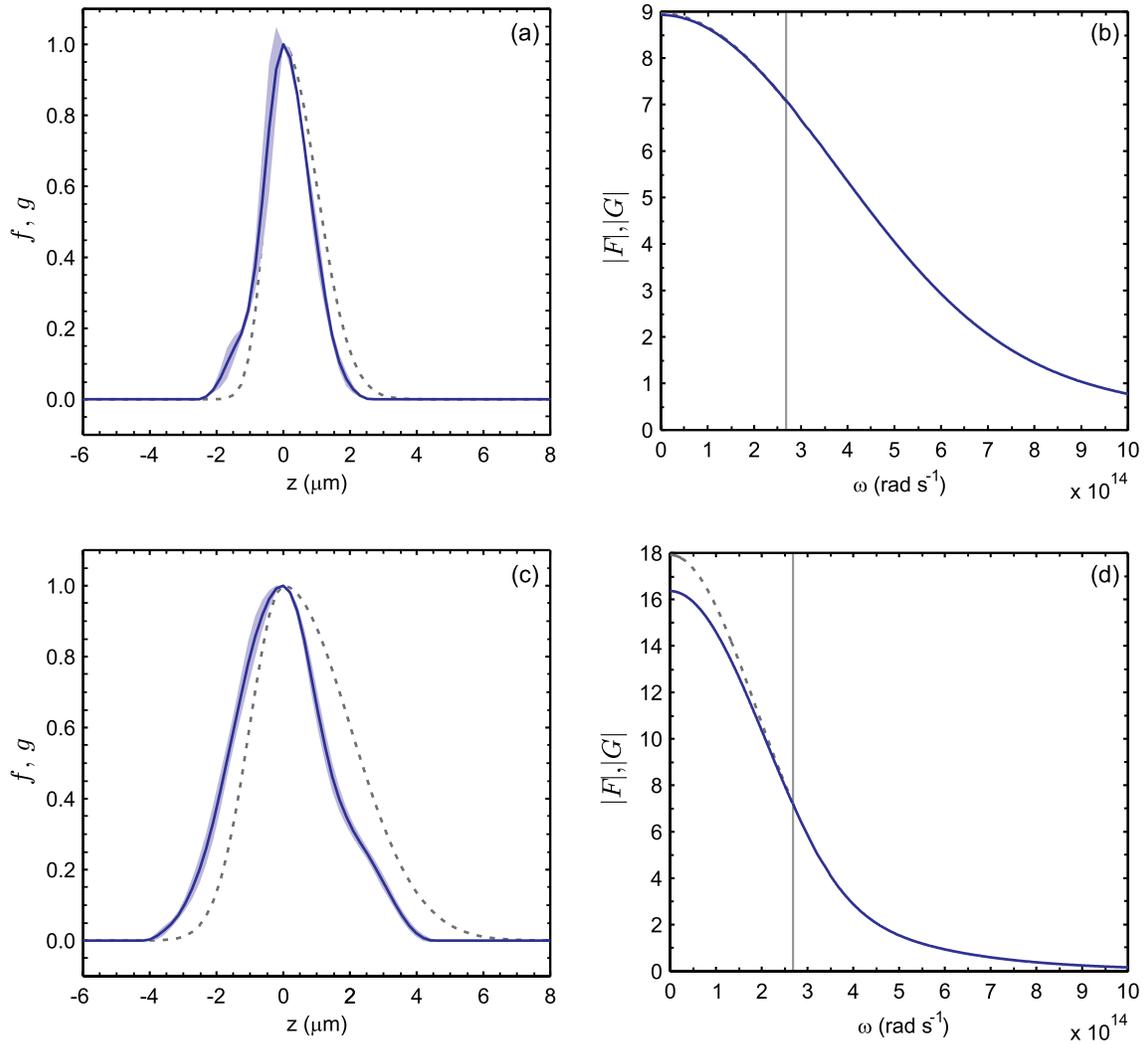


FIG. 3. Reconstruction of sample asymmetric bunch profiles. Initial profiles are shown as dashed gray lines in (a) and (c). The mean of the reconstructed profiles—from ten reconstructions with randomized initial data—are shown as blue lines. The light blue shading represents three- σ deviation from the mean. Profile (a) is a combination of half Gaussians of rms lengths $0.5 \mu\text{m}$ (to the left of $z = 0$) and $1 \mu\text{m}$ (to the right); for profile (c) the lengths are correspondingly $1 \mu\text{m}$ and $2 \mu\text{m}$. The spectral amplitudes corresponding to the original and reconstructed profiles are shown in (b) and (d). Spectral data is withheld from the phase-retrieval algorithm for $\omega < 2.7 \times 10^{14} \text{ rad s}^{-1}$, which is the only region where discrepancies between the original and reconstructed spectra occur.

data, at levels of 10% and 20% rms. For the reconstruction of regular Gaussian profiles, we found that noise introduced a random uncertainty in the reconstructed bunch lengths that is commensurate with, or less than, the applied noise level. Notably, applying such nonsystematic high-frequency noise in the Fourier domain results in the anomalous reconstruction of bunch features at long distances from the main bunch in the real-space domain—these can normally be disregarded on the basis of physical considerations. In order for measurement errors to introduce actual variations in the bunch shape—or the appearance of features close to the main bunch—they would need to be systematic, altering the actual shape of the recorded spectrum over a range of frequencies. Such errors could arise, for example, from a miscalibration of the detectors' spectral response.

In two dimensions, phase retrieval on experimentally recorded diffraction patterns (or other Fourier-domain data) can often lead to demonstrably unique real-space reconstructions. In general, this is not the case for phase retrieval in 1D [28], although functions with disconnected supports which satisfy a separation condition have been shown to almost always have a unique solution [29]. Hence, in general for the 1D case there may be profiles different from the reconstructed one that satisfy the same support conditions and have the same Fourier amplitudes. However, a number of indicators lead us to believe that the profiles reconstructed using the algorithm presented here are good approximations of the true inversions of the measured spectra. First, we find that reconstructions starting from different randomized seeds very often give similar reconstructed profiles. Second, these reconstructed

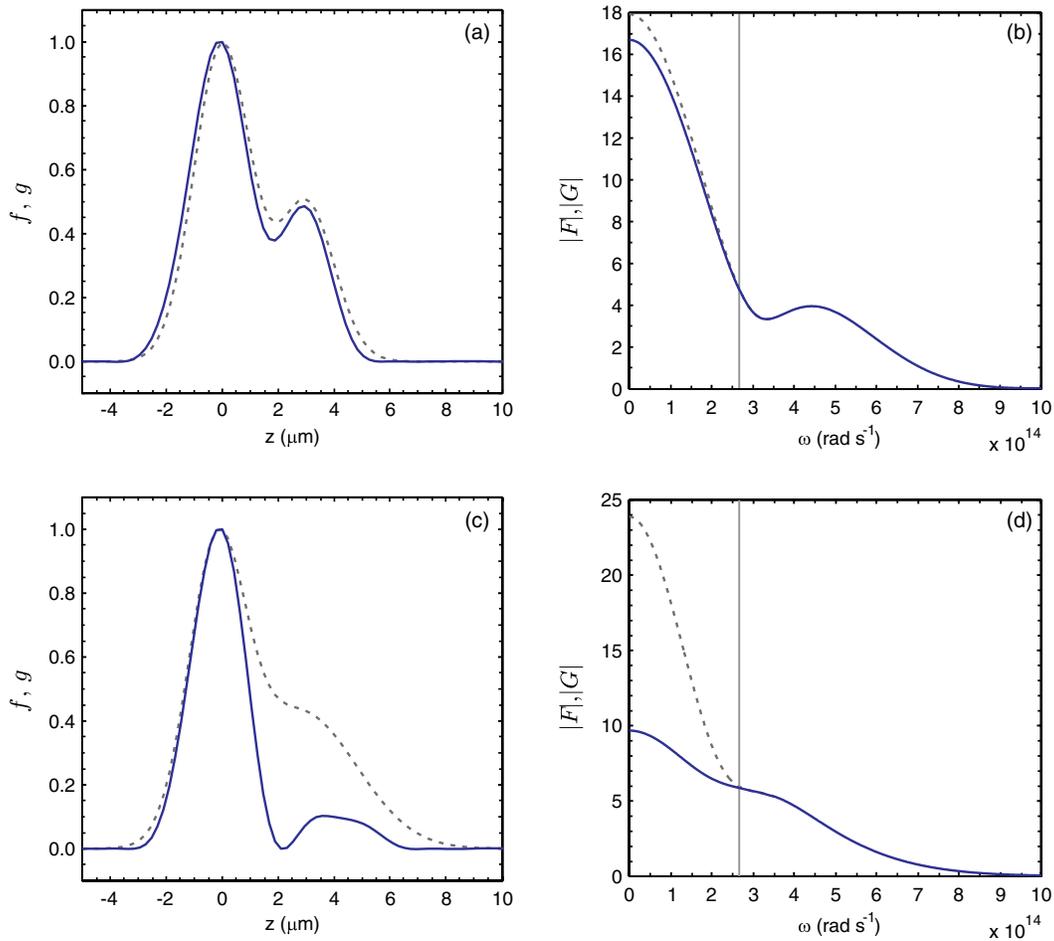


FIG. 4. Reconstruction of synthetic bunch profiles comprising two Gaussians, where the peak height of the smaller one is half that of the larger one. Initial profiles are shown as dashed gray lines in (a) and (c). The mean of the reconstructed profiles is shown as blue lines, and light blue shading represents three- σ deviation from the mean. In (a) both Gaussians have rms length $1 \mu\text{m}$, and their centers are separated by $3 \mu\text{m}$; in (c) the separation is the same, but the smaller Gaussian has an rms length of $2 \mu\text{m}$. The spectral amplitudes corresponding to the original and reconstructed profiles are shown in (b) and (d).

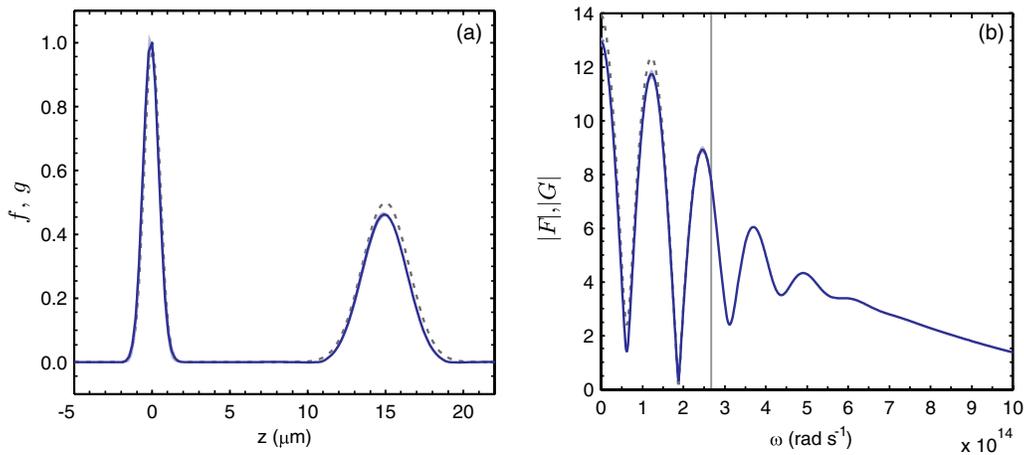


FIG. 5. Reconstruction of a synthetic bunch profile comprising two Gaussians. The initial profile is shown as dashed gray lines in (a). The mean of ten reconstructed profiles is shown as a blue line, and light blue shading represents three- σ deviation from the mean. The rms lengths of the two Gaussians are $0.5 \mu\text{m}$ and $1.5 \mu\text{m}$ for the larger and smaller one, respectively, and the separation is $15 \mu\text{m}$. The spectral amplitudes corresponding to the original and reconstructed profiles are shown in (b).

profiles are reasonable from a physical standpoint. Third, within a set parameter space tests with synthetic data give accurate reconstructed results. In conclusion, while it may be possible to find other reconstructions which are consistent with the measured spectra, they are unlikely to be realistic longitudinal bunch profiles.

V. MEASUREMENT OF CTR SPECTRA FROM LASER-ACCELERATED ELECTRON BUNCHES

Experimental measurements of the spatial and spectral characteristics of CTR emitted from laser-accelerated electrons were carried out at the Max-Planck-Institut für Quantenoptik (MPQ) in Garching, Germany. The ATLAS Ti:sapphire laser system at MPQ delivered an on-target laser energy of 1.5 ± 0.1 J within a FWHM pulse length of 28 ± 2 fs, focused to a waist of spot size of $w_0 = 18.7 \pm 1.2$ μm . This gave a peak intensity of $5.9 \pm 0.9 \times 10^{18}$ W cm^{-2} , corresponding to a peak normalized vector potential in vacuum of $a_0 = 1.66 \pm 0.13$.

The plasma target in which acceleration took place was a steady-state-flow gas cell, where the laser-plasma interaction length could be set from 2 to 14 mm. The pressure of the hydrogen gas in the cell was varied in the range 65–260 mbar, corresponding to plasma densities of $3.1\text{--}12.6 \times 10^{18}$ cm^{-3} . The experimental setup is shown schematically in Fig. 6.

Upon exiting the gas cell, the accelerated electron bunch passed through a pair of 20 μm -thick steel tapes located 50 mm behind the gas cell exit. The first tape served to block the laser beam, while the second tape was the source of transition radiation. Ballistic longitudinal bunch expansion over this short distance would be negligible. The transverse extent of the beam was accounted for in the spectrometer response function.

The transition radiation was reflected to one side by a pellicle placed directly behind the tape drive, while the electron bunch continued along the propagation axis of the

driving laser. We note that the tape and pellicle are both opaque to the detected radiation and hence only transition radiation from the rear of the second tape and the front of the pellicle was detected. These two sources of radiation may interfere, but since the distance between them is much smaller than the CTR formation length, the two sources will be close to being in phase over the range of angles and wavelengths detected. As such modulation of the spectrum or angular distribution of the detected transition radiation by interference effects will be negligible. Further, full Fourier optics source and transport calculations—to be published in detail elsewhere—show that at the detector the temporal and spatial overlap of these two sources of radiation is small, and hence that interference effects are weak.

The electron bunch was dispersed by a dipole magnet spectrometer, and the charge and energy spectrum of the bunch were recorded by means of detecting the fluorescence of an absolutely calibrated Lanex screen.

After reflection from the pellicle, the transition radiation was collimated by an off-axis paraboloid of effective focal length $f = 19$ cm and an f number of $f/3.75$, corresponding to an acceptance half angle of 133 mrad. Part of the collimated radiation propagated through a 1 mm-thick silicon wafer, and was reflected into a separate vacuum chamber that housed the THz spectrometer. This spectrometer was based on a design employed for CTR measurements at DESY, which is described in more detail by Delsim-Hashemi [30]. It employs three gratings; the zeroth order reflection from each is directed towards the next. The first grating works as a low pass filter to avoid ambiguities due to the overlap of different diffraction orders. The parts of the spectrum dispersed by the second and third grating are focused onto two pyroelectric detector arrays by custom-made gold mirrors. Two grating configurations were available, allowing measurements either in the range 1.7–7 μm or 5–25 μm . The raw detector signal was processed and digitized as described in Ref. [30].

The first silicon wafer mentioned above directed part of the radiation out of the vacuum chamber to a separate in-air spectrometry setup. There, a lens of 60 cm focal length was used to refocus the collimated transition radiation onto the entrance slits of an imaging spectrometer (Oriel MS260) and a near-infrared (NIR) spectrometer (Princeton Instruments OMA-C) which employed a 1024×256 silicon CCD chip and a 1024-element liquid nitrogen-cooled InGaAs diode array, respectively. The former spectrometer, hereafter referred to as the “visible” spectrometer, was used for detection in the spectral range 421–1096 nm, while the latter was used over three different ranges—each spanning approximately 700 nm—at wavelengths between 1033 and 2135 nm.

Since silicon has good reflectivity (≥ 0.3) over both the visible and near-infrared, but only transmits wavelengths longer than ~ 1050 nm [31], a second silicon wafer served

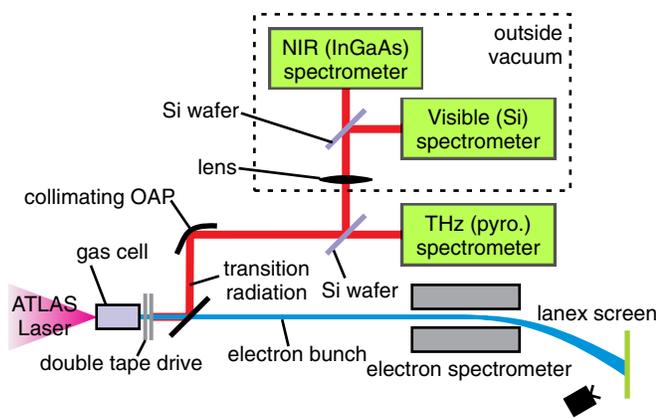


FIG. 6. Experimental setup for characterization of coherent transition radiation across the visible, near-infrared, and terahertz ranges of the electromagnetic spectrum. See text for details.

as the beam splitter between the visible and NIR spectrometers, which measured the reflected and transmitted radiation, respectively. By this token, a third silicon wafer (not pictured in Fig. 6) was placed directly at the entrance of the NIR spectrometer in order to eliminate stray laser radiation that could give an erroneous signal from higher-order reflection off the NIR spectrometer grating.

Through the use of an 1100 K tungsten blackbody source and an Ocean Optics HL-2000-CAL calibration source, the relative spectral responses of the visible and NIR spectrometers were established. A 50 mW helium-neon laser placed at the CTR source position was used with a chopper in order to obtain an absolute sensitivity calibration for the visible spectrometer, from which an absolute calibration of the NIR spectrometer could also be established.

Owing to lack of access to a well-characterized terahertz source, the relative calibration of the THz spectrometer was established according to the following procedure. The frequency response of the detection system was decomposed into (a) a response arising from the wavelength dependence of the grating reflectivity and the variation of the grating dispersion across the detector array; and (b) the response of each detector, accounting for the wavelength dependence of the general pyroelectric response, as well as variations between individual detectors. After accounting for the former, the latter was determined as follows. For a large number of shots where the laser and plasma parameters were expected to generate only a single electron bunch, as confirmed by the absence of obvious structure in the optical and NIR spectra, it was assumed that a single

electron bunch was indeed generated. The Bubblewrap algorithm was run for these shots, artificially restricting the support to only the main current peak. This yielded a smooth theoretical spectrum, which could be compared against the measured THz spectrum. The mean of the thus-obtained calibration factors was taken to obtain the relative calibration for each detector. After the relative response of the pyroelectric detectors had been obtained in this way, the measured spectra were corrected for variations in the pyroelectric response and in subsequent Bubblewrap runs the spectrum was constrained to the (corrected) measured spectrum.

Figure 7 illustrates how the overall relative calibration is made up from these two contributions to the THz frequency response. The absolute level of terahertz emission was determined separately, from the spectral region where the NIR and THz spectrometers overlapped.

VI. RECONSTRUCTION OF LONGITUDINAL BUNCH PROFILES FROM MEASURED CTR SPECTRA

The Bubblewrap algorithm was applied to 816 experimental shots for which the 1.7–7 μm grating configuration was used in the THz spectrometer and where good spectral data was available through the visible, near-infrared, and terahertz regions. A further 538 shots were taken using the 5–25 μm grating configuration, but the gap in the spectrum at 2–5 μm precluded reliable application of the phase-retrieval algorithm in those cases. Moreover, for this second data set, the lack of an absolute calibration or an overlap region with the calibrated NIR data made it difficult to determine the absolute level of the THz signal for these long-wavelength measurements. Nonetheless, measurements at the long-wavelength configuration confirm a smooth spectrum without structures for shots where a single electron bunch was expected, such as those used for calibration.

In this section we present the results of the bunch profile reconstruction from experimental CTR data for a small number of representative cases. We employed a grid of $N = 8192$ points for the reconstruction (hereafter the “reconstruction grid”), which in the Fourier domain spans frequencies from $\omega_{\min} = 0$ up to the highest measured frequency $\omega_{\max} = 2\pi c/(420 \text{ nm}) = 4.48 \times 10^{15} \text{ rad s}^{-1}$. This highest frequency also sets the data point spacing in real space to be $\Delta z = 210 \text{ nm}$, since by Nyquist’s theorem $\omega_{\max} = 2\pi c/(2\Delta z)$. The number of data points then sets the total span of the reconstruction window to $N\Delta z \approx 1.72 \text{ mm}$. Since the total length of the electron bunch—including secondary features—is unlikely to exceed 50 μm , the reconstruction is largely insensitive to N , as long as $N \gtrsim 500$. This was confirmed with synthetic data, showing that equivalent reconstructions are obtained for grids of 2048 and 8192 points.

Since the data from the visible and NIR spectrometers has a resolution that is higher than (or comparable to) the

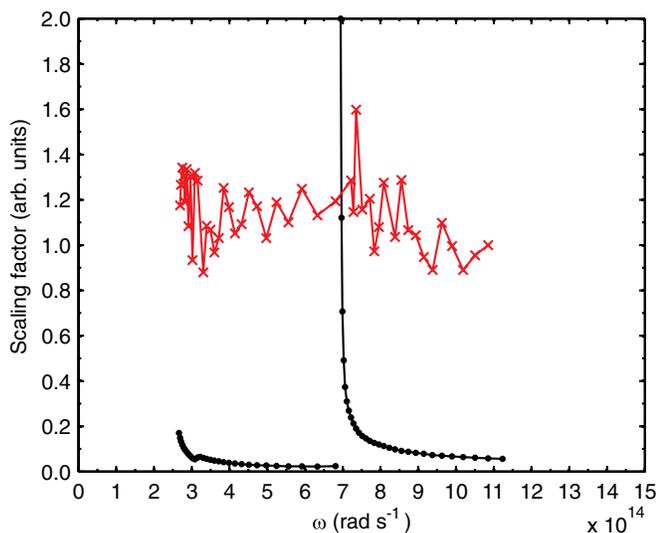


FIG. 7. Contributions to the relative pyroelectric detector calibration due to (a) grating reflectivity and dispersion (black dots); and (b) pyroelectric detector response (red crosses). The calibration factors plotted represent the inverse of system/detector sensitivity. Each data point represents a single pyroelectric detector. The significant variation between individual detector responses illustrates the need for reliable calibration.

resolution of the reconstruction grid in Fourier space, linear interpolation was used to populate the initial candidate spectrum function, G_0 , with the amplitude data (i.e., the square roots of the measured intensities). On the other hand, the terahertz data has a relatively low resolution, with fewer data points than the reconstruction grid. In this case, G_0 was populated by taking each of the terahertz data points and assigning its value to the nearest corresponding point on the reconstruction grid. The remaining points on the grid were left unassigned, and their values were free to vary during the phase-retrieval process.

For each measured spectrum the Bubblewrap algorithm was run 10 times using different random seeds, allowing the stability of the reconstruction to be evaluated. This also allowed a representative profile to be selected, as follows. Starting with all ten reconstructed bunch shapes, the profile that was most different from the others—in a least squares sense—was removed from the set. This elimination was repeated until only three profiles were left. Of these, the profile most similar to the other two was selected as the most representative of the set. This process appeared to yield valid results for the majority of the data. Even in cases where the reconstruction produced diverging results for some initial seeds, eliminating the outliers ultimately leads to the selection of a representative profile.

Figures 8–10 show the reconstructed profiles and measured and reconstructed CTR spectra for three representative single-shot measurements, together with a summary of acceleration and the deduced electron bunch parameters. The recorded images of the Lanex screen after the electron spectrometer are also shown for reference. Note that panels (ii) and (iii) in each figure show the results of all ten separate reconstructions. Despite the random seeding, the ultimate reconstructed profiles demonstrate a striking degree of convergence. The presented figures illustrate the general shape of the reconstructed profiles—occasionally with secondary bunches, the interference of which causes the observed oscillations in the CTR spectrum.

In order to establish the validity of the reconstructed profiles we may compare the reconstructed bunch lengths with the condition $\sigma_z \lesssim \lambda_c/2\pi$. The full width at half maximum (FWHM) lengths of the reconstructed bunches are shown in Figs. 8–10. Given that the FWHM of a Gaussian profile of rms length σ_z is $2\sqrt{2\ln 2}\sigma_z$, we see that the deduced bunch lengths are well within the range for which the reconstruction is expected to be reliable.

The validity of the bunch reconstructions can also be checked by comparing the total energy contained in the measured CTR spectra, against that expected based on the reconstructed bunch profile. To do this, we can take as a starting point the formula for total energy emitted by a single electron, $dU_1/d\omega = e^2/(2\pi^2\epsilon_0 c) \ln\gamma$, which is the integral of the Ginzburg-Frank formula over all angles in the limit $\gamma \gg 1$ (after Ref. [32]), where γ is the Lorentz factor of the electron. Taking into account the total number

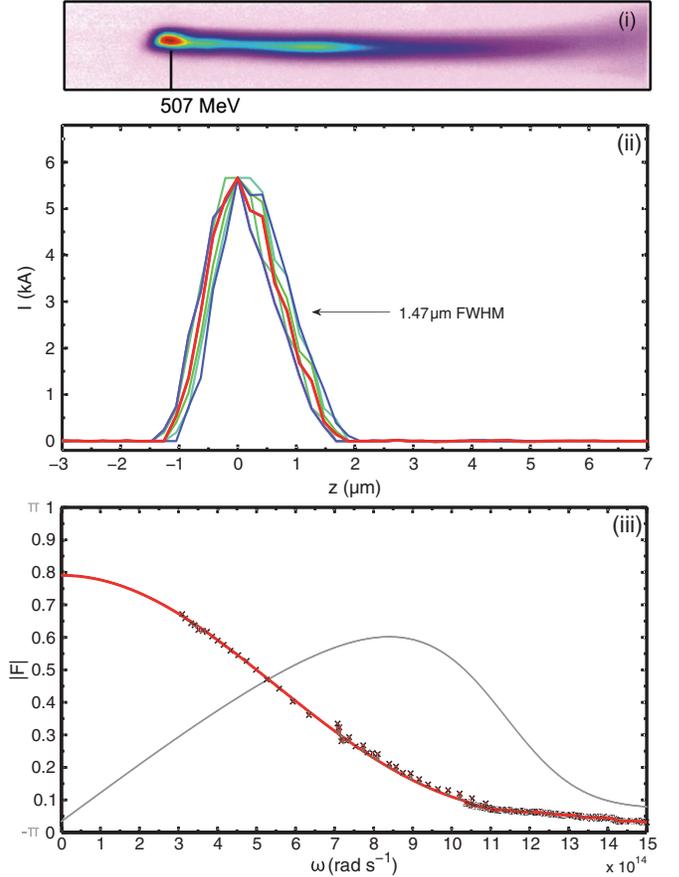


FIG. 8. Reconstruction of longitudinal bunch profile from experimentally measured CTR spectrum: (i) image of Lanex screen at the exit of the electron spectrometer; (ii) ten reconstructed profiles, the most representative of which is colored red; (iii) measured spectrum (thick black line, black crosses) and reconstructed spectra (colored lines), including retrieved phase of most representative profile (gray dots). Data for Shot 1930: $P = 83$ mbar ($n_e = 4.0 \times 10^{18} \text{ cm}^{-3}$); $L_{\text{cell}} = 6$ mm; $E_{\text{peak}} = 507 \pm 92$ MeV (FWHM); $Q_{\text{peak}} = 8.7$ pC; $Q_{\text{total}} = 34.6$ pC.

of electrons, N_e , and the frequency-dependent longitudinal coherence factor, $\tilde{\rho}_{\parallel}(\omega)$, the total coherent spectrum is given by

$$\frac{dU}{d\omega} = N_e^2 |\tilde{\rho}_{\parallel}(\omega)|^2 \frac{e^2}{2\pi^2\epsilon_0 c} \langle \ln\gamma \rangle, \quad (7)$$

where $\langle \ln\gamma \rangle = \frac{\int dE \ln\gamma(E) (dQ/dE)}{\int dE (dQ/dE)}$ is the weighted average of $\ln\gamma$ over the electron energy spectrum, dQ/dE . Because of the small variation in $\ln\gamma$ ($\approx 6-7$) over the range of observed electron energies, the use of this weighted average to represent the scaling with γ —instead of separately considering the emission from parts of the electron bunch with different energies—should not lead to a significant discrepancy.

The total number of electrons in the bunch can be found by calibrating the response of the Lanex screen used to

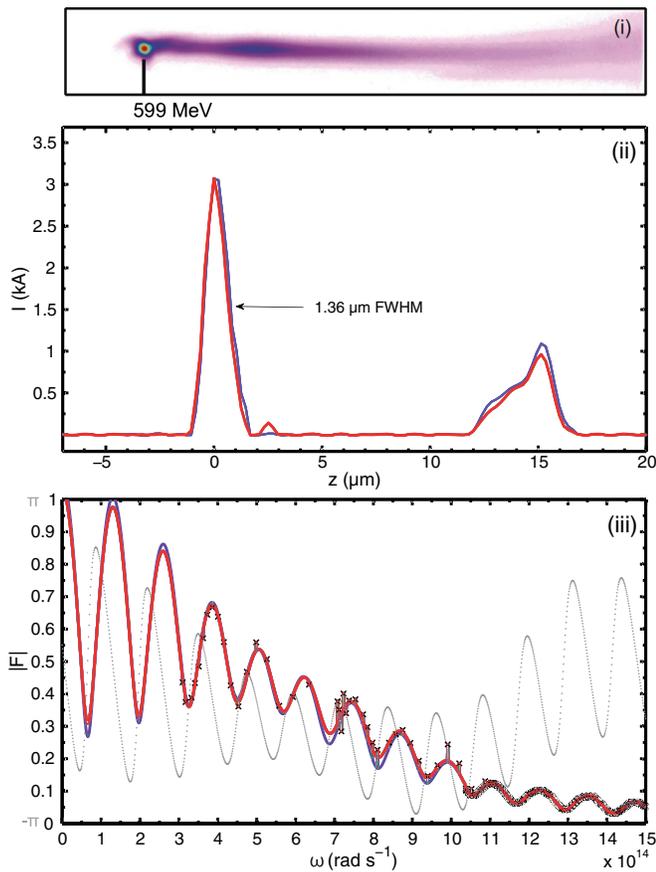


FIG. 9. Reconstruction of longitudinal bunch profile from experimentally measured CTR spectrum: (i) image of Lanex screen at the exit of the electron spectrometer; (ii) ten reconstructed profiles, the most representative of which is colored red; (iii) measured spectrum (thick black line, black crosses) and reconstructed spectra (colored lines), including retrieved phase of most representative profile (gray dots). Data for Shot 2168: $P = 82$ mbar ($n_e = 4.0 \times 10^{18} \text{ cm}^{-3}$); $L_{\text{cell}} = 14$ mm; $E_{\text{peak}} = 599 \pm 47$ MeV; $Q_{\text{peak}} = 3.8$ pC; $Q_{\text{total}} = 22.5$ pC.

detect the dispersed electron energy spectrum [33] and integrating over the measured energy spectrum. Now, the total number of electrons in the bunch is also given by integrating over the longitudinal bunch profile: $N_e = N_e \int \rho_{\parallel}(z) dz = N_e \tilde{\rho}(0)$, where the right-hand side follows from considering the Fourier transform of $\rho_{\parallel}(z)$ at $\omega = 0$. Normalizing the measured spectrum $dU/d\omega$ by $N_e^2 (e^2/2\pi^2 \epsilon_0 c) \langle \ln \gamma \rangle$ yields $|\tilde{\rho}(\omega)|$ and hence $|\tilde{\rho}(0)|$. If it is found that $|\tilde{\rho}(0)| = 1$, then the total bunch charge deduced from the reconstructed bunch profile agrees with that measured independently by the electron spectrometer. Within a reasonable margin of error, we observe this to be true in most cases, and certainly for those shown in Figs. 8–10.

It is possible that not all electrons that emit transition radiation reach the electron spectrometer. Additionally, the absolute calibration of the THz detectors is only obtained

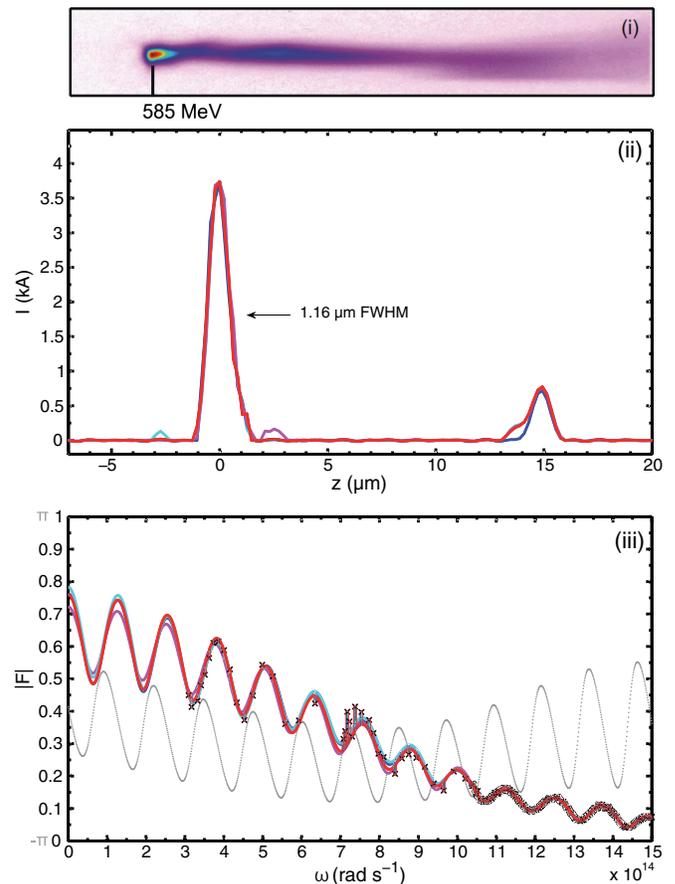


FIG. 10. Reconstruction of longitudinal bunch profile from experimentally measured CTR spectrum: (i) image of Lanex screen at the exit of the electron spectrometer; (ii) ten reconstructed profiles, the most representative of which is colored red; (iii) measured spectrum (thick black line, black crosses) and reconstructed spectra (colored lines), including retrieved phase of most representative profile (gray dots). Data for Shot 2147: $P = 82$ mbar ($n_e = 4.0 \times 10^{18} \text{ cm}^{-3}$); $L_{\text{cell}} = 13$ mm; $E_{\text{peak}} = 585 \pm 48$ MeV; $Q_{\text{peak}} = 2.8$ pC; $Q_{\text{total}} = 21.0$ pC.

by cross calibration with the NIR spectrometer at a small number of data points around $2 \mu\text{m}$, and may be unreliable. Therefore, the above verification of reconstruction validity should only be used as a rough guide. Nonetheless, we find $\tilde{\rho}$ to be close to 1 for a majority of the reconstructed spectra. This suggests that most of the emitted radiation was collected and the absolute calibration of the spectrometers is approximately accurate.

As discussed in the Introduction, other methods exist for deducing the longitudinal bunch profile from the measured CTR spectrum. A rudimentary method of estimating the bunch length involves assuming that the longitudinal bunch profile is Gaussian and accordingly fitting the measured spectrum to a Gaussian in Fourier space. Applying this method to the available experimental data, we find that the FWHM lengths of the bunch profiles deduced by the Bubblewrap algorithm are almost universally in the range

1.1 ± 0.1 times the FWHM lengths obtained by Gaussian fitting. Such a small discrepancy gives further credibility to the bunch profiles obtained by phase retrieval. The existence of the discrepancy itself highlights the fact that the assumption of a Gaussian profile is likely to limit the accuracy of bunch length estimation by simple fitting, by, for example, not accounting for additional bunches or asymmetry.

VII. CONCLUSIONS

We have considered the application of phase-retrieval algorithms to reconstruction of the longitudinal bunch profile of an electron bunch from measurements of the coherent transition radiation spectrum it produces. Several existing algorithms were investigated, though no single one alone was found able to perform the reconstruction reliably for a sufficiently wide range of potential scenarios. Based on a combination of these algorithms, a new algorithm—the Bubblewrap algorithm—was developed for this application. Tests with synthetic data showed that it was able to reconstruct the longitudinal profile of single and double electron bunches, provided that the CTR spectrum was known over a sufficiently wide range. The algorithm was also applied to recently measured CTR spectra generated by laser-accelerated electron bunches. The retrieved longitudinal bunch profiles were found to be realistic, and consistent with estimates of the bunch length obtained by assuming a Gaussian profile. Retrieval of the bunch profile takes a few seconds on a desktop personal computer without code optimization, suggesting the possibility of implementing this as an online diagnostic for accelerators operating at higher repetition rates.

We note that charged particle beams can generate radiation by a wide range of phenomena—including coherent transition radiation, coherent diffraction radiation, and Smith-Purcell radiation—but direct reconstruction of the longitudinal bunch shape is not possible since the spectral phase is not known. The analysis of these spectra by phase-retrieval algorithms offers one solution to this problem, and has several advantages over other methods. The algorithms are tolerant to gaps in the measured spectra, including missing information about very low frequencies. They also provide much more information about the bunch profile than simpler methods, such as fitting the spectrum to that calculated for a parametrized bunch profile. For example, as long as the spectrum is measured over a sufficiently wide range, multiple bunches and bunch profile asymmetry can all be detected. The additional information provided by phase-retrieval algorithms, such as Bubblewrap, is likely therefore to be useful in illuminating the physical processes occurring during the bunch generation.

ACKNOWLEDGMENTS

We would like to acknowledge very useful discussions on this topic with Armin Reichold, Riccardo Bartolini,

George Doucas, and Richard Tovey. We would also like to acknowledge the collaboration of Bernhard Schmidt's group at DESY on the design and construction of the THz spectrometer. This work was supported by the Engineering and Physical Sciences Research Council (Grant No. EP/H011145/1), the Leverhulme Trust (Grant No. F/08 776/G), the Max-Planck Society, the DFG (MAP and TR-18 funding schemes), and the Euratom-IPP association.

- [1] V.L. Ginzburg and V.N. Tsytovich, *Transition Radiation and Transition Scattering* (Hilger, Bristol, UK, 1990).
- [2] W.P. Leemans, C.G.R. Geddes, J. Faure, C. Tóth, J. van Tilborg, C.B. Schroeder, E. Esarey, G. Fubiani, D. Auerbach, B. Marcellis, M.A. Carnahan, R.A. Kaindl, J. Byrd, and M.C. Martin, *Phys. Rev. Lett.* **91**, 074802 (2003).
- [3] J. Faure, C. Rechatin, A. Norlin, A. Lifschitz, Y. Glinec, and V. Malka, *Nature (London)* **444**, 737 (2006).
- [4] Y. Glinec, J. Faure, A. Norlin, A. Pukhov, and V. Malka, *Phys. Rev. Lett.* **98**, 194801 (2007).
- [5] G. Berden, S.P. Jamison, A.M. MacLeod, W.A. Gillespie, B. Redlich, and A.F.G. van der Meer, *Phys. Rev. Lett.* **93**, 114802 (2004).
- [6] G. Berden, W.A. Gillespie, S.P. Jamison, E.A. Knabbe, A.M. MacLeod, A.F.G. van der Meer, P.J. Phillips, H. Schlarb, B. Schmidt, P. Schmüser, and B. Steffen, *Phys. Rev. Lett.* **99**, 164801 (2007).
- [7] F.J. Grüner, S. Becker, U. Schramm, T. Eichner, M. Fuchs, R. Weingartner, D. Habs, J. Meyer-Ter-Vehn, M. Geissler, M. Ferrario, L. Serafini, B. van der Geer, H. Backe, W. Lauth, and S. Reiche, *Appl. Phys. B* **86**, 431 (2007).
- [8] R.B. Fiorito and D.W. Rule, *AIP Conf. Proc.* **319**, 21 (1994).
- [9] G.P. Le Sage, T.E. Cowan, R.B. Fiorito, and D.W. Rule, *Phys. Rev. ST Accel. Beams* **2**, 122802 (1999).
- [10] R. Lai, U. Happek, and A.J. Sievers, *Phys. Rev. E* **50**, R4294 (1994).
- [11] A. Murokh, J.B. Rosenzweig, M. Hogan, H. Suk, G. Travish, and U. Happek, *Nucl. Instrum. Methods Phys. Res., Sect. A* **410**, 452 (1998).
- [12] S. Wesch, B. Schmidt, C. Behrens, H. Delsim-Hashemi, and P. Schmüser, *Nucl. Instrum. Methods Phys. Res., Sect. A* **665**, 40 (2011).
- [13] O. Lundh, J. Lim, C. Rechatin, L. Ammoura, A. Ben-Ismaïl, X. Davoine, G. Gallot, J.P. Goddet, E. Lefebvre, V. Malka, and J. Faure, *Nat. Phys.* **7**, 219 (2011).
- [14] R. Lai and A.J. Sievers, *Phys. Rev. E* **52**, 4576 (1995).
- [15] R. Lai and A.J. Sievers, *Nucl. Instrum. Methods Phys. Res., Sect. A* **397**, 221 (1997).
- [16] R. Lai and A.J. Sievers, *Phys. Rev. E* **50**, R3342 (1994).
- [17] G. Schneider, M.G. Billing, R. Lai, and A.J. Sievers, *Phys. Rev. E* **56**, R3780 (1997).
- [18] D. Mihalcea, C.L. Bohn, U. Happek, and P. Piot, *Phys. Rev. ST Accel. Beams* **9**, 082801 (2006).
- [19] J.R. Fienup, *Appl. Opt.* **21**, 2758 (1982).
- [20] S. Marchesini, H. He, H.N. Chapman, S.P. Hau-Riege, A. Noy, M.R. Howells, U. Weierstall, and J.C.H. Spence, *Phys. Rev. B* **68**, 140101 (2003).

- [21] O. Shapira, A.F. Abouraddy, J.D. Joannopoulos, and Y. Fink, *Phys. Rev. Lett.* **94**, 143902 (2005).
- [22] D. Sayre, *Acta Crystallogr.* **5**, 843 (1952).
- [23] R.W. Gerchberg and W.O. Saxton, *Optik* **35**, 237 (1972).
- [24] V. Elser, *J. Opt. Soc. Am. A* **20**, 40 (2003).
- [25] S. Marchesini, S. Boutet, A.E. Sakdinawat, M.J. Bogan, S. Bajt, A. Barty, H.N. Chapman, M. Frank, S.P. Hau-Riege, A. Szoke, C. Cui, D.A. Shapiro, M.R. Howells, J.C. Spence, J.W. Shaevitz, J.Y. Lee, J. Hajdu, and M.M. Seibert, *Nat. Photonics* **2**, 560 (2008).
- [26] V. Elser, I. Rankenburg, and P. Thibault, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 418 (2007).
- [27] H.N. Chapman, A. Barty, S. Marchesini, A. Noy, S.P. Hau-Riege, C. Cui, M.R. Howells, R. Rosen, H. He, J.C.H. Spence, U. Weierstall, T. Beetz, C. Jacobsen, and D. Shapiro, *J. Opt. Soc. Am. A* **23**, 1179 (2006).
- [28] A. Walther, *Opt. Acta* **10**, 41 (1963).
- [29] T.R. Crimmins and J.R. Fienup, *J. Opt. Soc. Am.* **73**, 218 (1983).
- [30] H. Delsim-Hashemi, Ph.D. thesis, DESY, 2008.
- [31] M.A. Green and M.J. Keevers, *Prog. Photovoltaics* **3**, 189 (1995).
- [32] C.B. Schroeder, E. Esarey, J. van Tilborg, and W.P. Leemans, *Phys. Rev. E* **69**, 016501 (2004).
- [33] A. Buck, K. Zeil, A. Popp, K. Schmid, A. Jochmann, S.D. Kraft, B. Hidding, T. Kudyakov, C.M.S. Sears, L. Veisz, S. Karsch, J. Pawelke, R. Sauerbrey, T. Cowan, F. Krausz, and U. Schramm, *Rev. Sci. Instrum.* **81**, 033301 (2010).