Super-Resolving a Frequency Band

Ruiming Guo and Thierry Blu

This note introduces a simple formula that provides the exact frequency of a pure sinusoid from just two samples of its Discrete-Time Fourier Transform (DTFT). Even when the signal is not a pure sinusoid, this formula still works in a very good approximation (optimally after a single refinement), paving the way for high-resolution frequency tracking of fastly-varying signals, or simply improving the frequency resolution of the peaks of a Discrete Fourier Transform (DFT).

Single Frequency Estimation

We learn (and teach!) in college that the frequency content of a signal is encrypted in its Fourier transform, and that the main frequency modes of a sampled signal are the peaks of its DFT. We also learn that the accuracy of these frequency values is limited by the inverse of the time range of the signal (Heisenberg uncertainty), which correlates nicely with the inherent frequency resolution of a DFT: $2\pi/N$, if N is the number of samples of the signal. This knowledge is so deep-rooted that it is hard to reconcile with the fact that the frequencies of a signal made of a sum of K complex exponentials can be recovered *exactly* from as few as $2K$ samples, using a two-century-old method due to Gaspard de Prony. This apparent contradiction is resolved by recognizing that Heisenberg uncertainty relies on a much weaker signal assumption (basically, that its time and frequency uncertainties are finite) than the sum-of-exponentials model. And it is only when this model is inexact that the estimated frequencies may be inaccurate, with their uncertainty now given by Cramér-Rao lower bound which assumes unbiased estimators and known noise statistics. The contrast between the Fourier approach (analytic, intuitive, but Heisenberg-limited), and Prony's method (algebraic, black-box, but exact) has made it difficult to envision higher frequency resolution that would rely on the DFT. Yet, given that the DFT coefficients are just samples of a very smooth function, the DTFT, analytic considerations suggest that this function can be approximated locally by a quadratic polynomial, leading to an estimate of its off-grid peak location: as few as three samples around the max of the DFT already provide a very good estimate of this frequency [1], [2].

The Trick

The main motivation for this note is to put forward a formula which provides the frequency value of the maximum of the DTFT, from just two DTFT coefficients of a single-frequency signal: see detailed setting in Box 1. Not only is this formula exact, but it is also very robust to inaccuracies of the model, as we shall see later. More precisely, if we assume that the unknown frequency ω_0 of the signal x_n lies inside an "uncertainty band" $[\omega_1, \omega_2]$ where $\omega_2 - \omega_1$ is an integer multiple of $2\pi/N$, this formula specifies how

Correspondence to: Thierry Blu (thierry.blu@m4x.org), Department of Electronic Engineering, the Chinese University of Hong Kong, Sha Tin N.T., Hong Kong SAR of China

BOX 1 Notations

- Single-frequency signal: $x_n = a_0 \mathrm{e}^{jn\omega_0}$, $n = 0, 1, \dots N-1$
- Uncertainty band: $\omega_0 \in [\omega_1, \omega_2]$, with $\omega_2 - \omega_1 = \text{integer} \times 2\pi/N$
- Discrete Time Fourier Transform (DTFT): $X(\omega) = \sum_{n=0}^{N-1} x_n \mathrm{e}^{-jn\omega}$
- Discrete Fourier Transform (DFT): $X(2\pi k/N)$, where $k = 0, 1, \dots, N-1$
- Peak of the DFT: $k_0 = \arg\max_k |X(2\pi k/N)| \rightsquigarrow \omega_{\mathrm{DFT}} = 2\pi k_0/N$

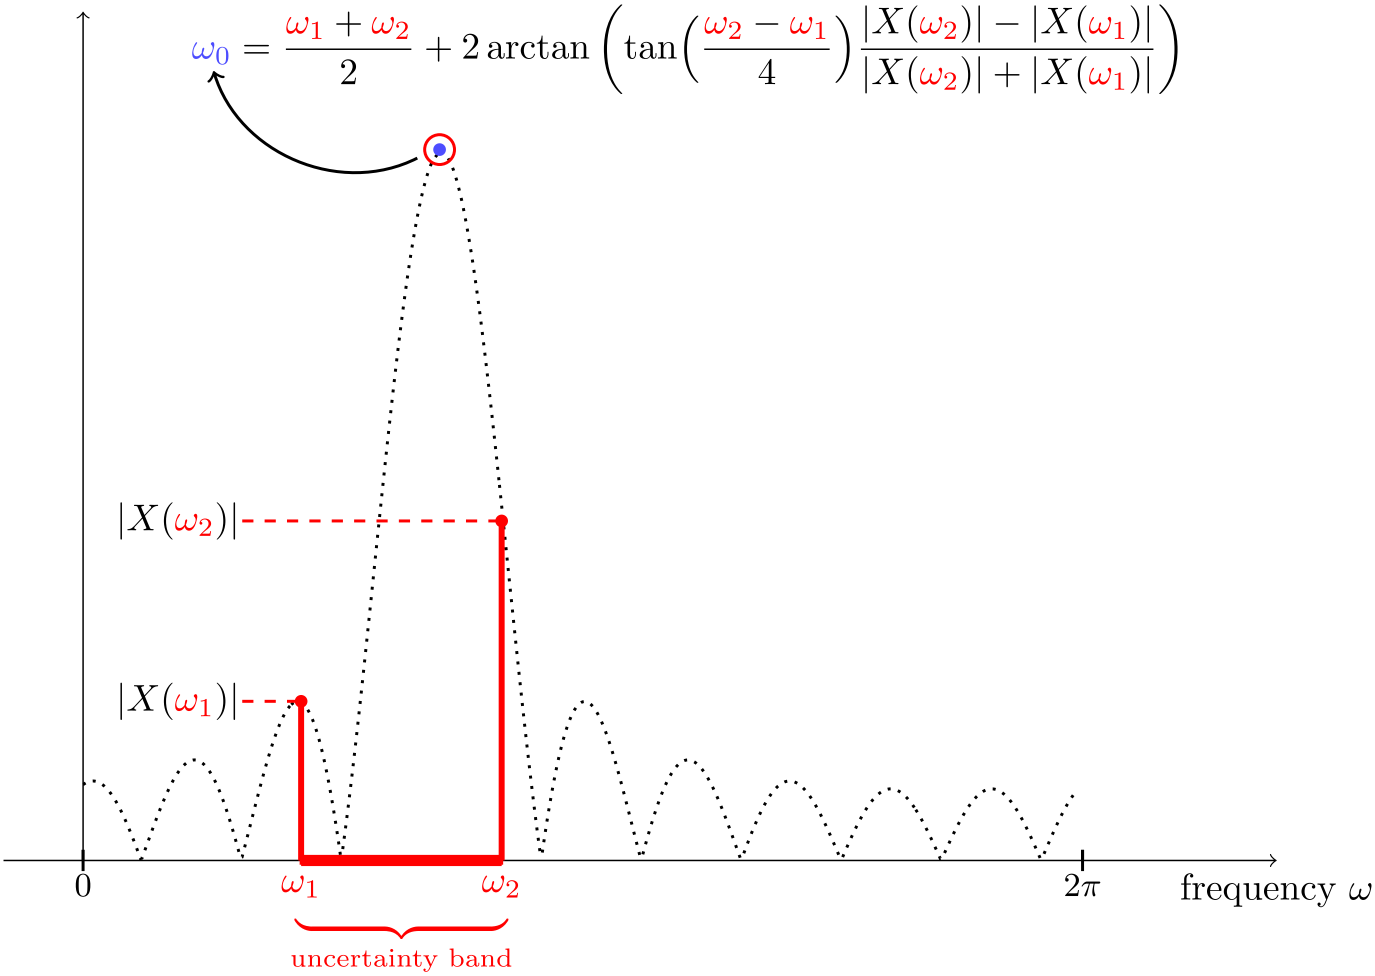

Fig. 1. The frequency ω_0 (blue) of a single-frequency signal is obtained from the DTFT values at the endpoints (red) of a frequency interval known to contain ω_0 (assumption: $\omega_2 - \omega_1 = \text{integer} \times 2\pi/N$). Dotted line: a full period of the DTFT $X(\omega)$ of the signal.

the amplitudes of the DTFT of x_n at ω_1 and ω_2 should be combined so as to recover ω_0 (see visualization in Fig. 1):

$$\omega_0 = \frac{\omega_1 + \omega_2}{2} + 2\arctan\left(\tan\left(\frac{\omega_2 - \omega_1}{4}\right)\frac{|X(\omega_2)| - |X(\omega_1)|}{|X(\omega_2)| + |X(\omega_1)|}\right). \tag{1}$$

A proof is provided in Box 2, requiring only elementary EEE-math knowledge. This formula becomes even simpler when the uncertainty bandwidth is small (i.e., $\omega_2 - \omega_1 \ll \pi$):

$$\omega_0 \approx \omega_1 \times \frac{|X(\omega_1)|}{|X(\omega_1)| + |X(\omega_2)|} + \omega_2 \times \frac{|X(\omega_2)|}{|X(\omega_1)| + |X(\omega_2)|},$$

which is what intuition would suggest: a weighting of the end frequencies based on the relative magnitude of their DTFT.

In practice, this formula is most useful when the uncertainty band is smallest; i.e., $\omega_2 - \omega_1 = 2\pi/N$. Then, a straightforward procedure for estimating a single frequency from N uniform signal samples is to:

1) determine ω_{DFT}, the frequency of the peak of the signal DFT;
2) apply (1) with $\omega_1 = \omega_{\text{DFT}} - \pi/N$ and $\omega_2 = \omega_{\text{DFT}} + \pi/N$.

This works because, for a single-frequency signal, the peak of the DTFT is always within $\pm\pi/N$ of the maximum of the DFT. But it would also be possible to bypass the computation of the full DFT if a rough estimate of ω_0 were available; e.g., when the frequency of the signal is continuously changing (tracking between successive signal windows like in radar applications), or when the frequency is a priori known up to some perturbation (physical resonance experiments, laser-based optical measurements, etc.).

Robustness to Inaccuracies

When the single-exponential model is not exact, the formula (1) is still a very robust estimator of its frequency. This is particularly so when the uncertainty bandwidth $\omega_2 - \omega_1$ is reduced to $2\pi/N$, an assumption that we will make from now on. Indeed, consider a noise model $x_n = a_0 \mathrm{e}^{jn\omega_0} + b_n$ where the complex-valued samples b_n are independent realizations of a Gaussian random variable with variance σ^2—"additive white Gaussian noise" assumption. When the number of samples N is large enough, a linearization of (1) makes it possible to calculate (tedious, not shown here) the standard deviation $\Delta\omega$ of the frequency estimation error. In particular, in the case where ω_0 is at the center of the interval $[\omega_1, \omega_2]$, this error behaves according to

$$\Delta\omega = \frac{\sqrt{2}\,\pi^2/4}{N\sqrt{N}\,\text{SNR}} = \frac{\pi^2}{4\sqrt{6}}\,\Delta\omega_{\text{CR}} \tag{2}$$

where $\text{SNR} = |a_0|/\sigma$ and where $\Delta\omega_{\text{CR}}$ is the Cramér-Rao lower bound of the problem (see [3] for a calculation). Obviously, $\Delta\omega$ is very close to $\Delta\omega_{\text{CR}}$, within less than 1%. When ω_0 is closer to the extremities of the interval $[\omega_1, \omega_2]$, $\Delta\omega$ deviates from the Cramér-Rao lower bound by up to 80%, still a very low error in absolute terms. In fact, a simple refinement of the trick as depicted in Fig. 2 shows how to attain this lower bound, outlining the near-optimality of this procedure: no other unbiased single-frequency estimation algorithm would be able to improve this performance by more than 1%.

This is confirmed in Fig. 3 by simulations that consist of one million tests, where: 1) the number of samples, N, is random (uniform) between 10 and 1000; 2) the frequency ω_0 is random (uniform) between $-\pi$ and π; 3) $a_0 = 1$; 4) the standard deviation σ of the noise is such that the SNR is random (uniform) between $-5\,\text{dB}$ and $40\,\text{dB}$; 5) the noise b_n is drawn from an iid statistics (Gaussian). In each test, the uncertainty band before refinement is set by $\omega_1 = \omega_{\text{DFT}} - \pi/N$ and $\omega_2 = \omega_{\text{DFT}} + \pi/N$. For comparison purposes, Fig. 3 also shows the distribution of errors of Jacobsen's estimator [1], [2], which is based on three consecutive DFT coefficients around ω_{DFT}. The better performance of our formula is likely due to the higher SNR enjoyed by the two DFT coefficients around the maximum of the DTFT, in comparison to the three DFT coefficients used by Jacobsen's formula, one of which has a significantly lower SNR, being further away from the DTFT peak by more than $2\pi/N$. A somewhat milder difference is also

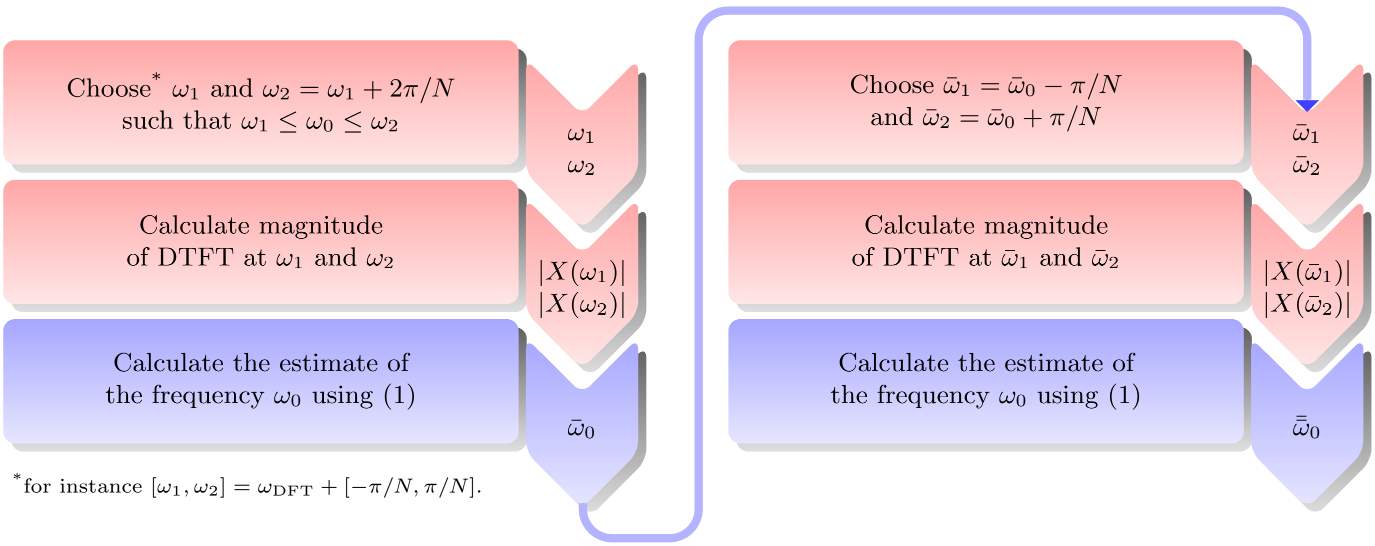

Fig. 2. Refined trick: a double application of (1) achieves a quality that is equivalent to the best unbiased single-frequency estimation algorithm.

that (1) assumes an exact single-frequency model, whereas Jacobsen's formula is but a local quadratic approximation.

Beyond just noise, when the single-frequency model is rendered inaccurate due to, e.g., quantization, sample windowing, or addition of other sinusoidal/polynomial terms, the frequency estimation error of (1) is controlled by ε, the maximum error of the magnitude of the DTFT at the frequencies ω_1 and ω_2, according to (proof in Appendix A)

$$|\bar{\omega}_0 - \omega_0| \le \underbrace{\frac{4\varepsilon \tan\left(\dfrac{\pi}{2N}\right)}{N|a_0|}}_{\approx \frac{2\pi\varepsilon}{N^2|a_0|} \text{ for large } N}. \tag{3}$$

where $|a_0|$ is the amplitude of the single-frequency "ground-truth" signal.

Note that, because it is valid for every single "noise" instance, this bound is of a very different nature than the statistical result (2)—an average over infinitely many additive white Gaussian noise realizations. Interestingly, inaccuracies of the model outside the uncertainty band do not contribute to the estimation error, which suggests that (3) can be used to predict the accuracy of a multiple frequency estimation problem that uses the single-frequency trick.

Multiple Frequencies

This formula can easily be used in a multiple frequency scenario, provided that the frequencies to estimate are sufficiently separated. A straightforward procedure consists in, first locating the isolated peaks of the DFT of the signal (e.g., using Matlab's `findpeaks` function), then applying (1) to refine each frequency individually: see an example in Fig. 4. The estimation error of each frequency can be quantified by using the bound (3) where, in the absence of other noise, the data inaccuracy ε in the neighborhood of that frequency is essentially caused by the tail of the DTFT of the other frequencies.

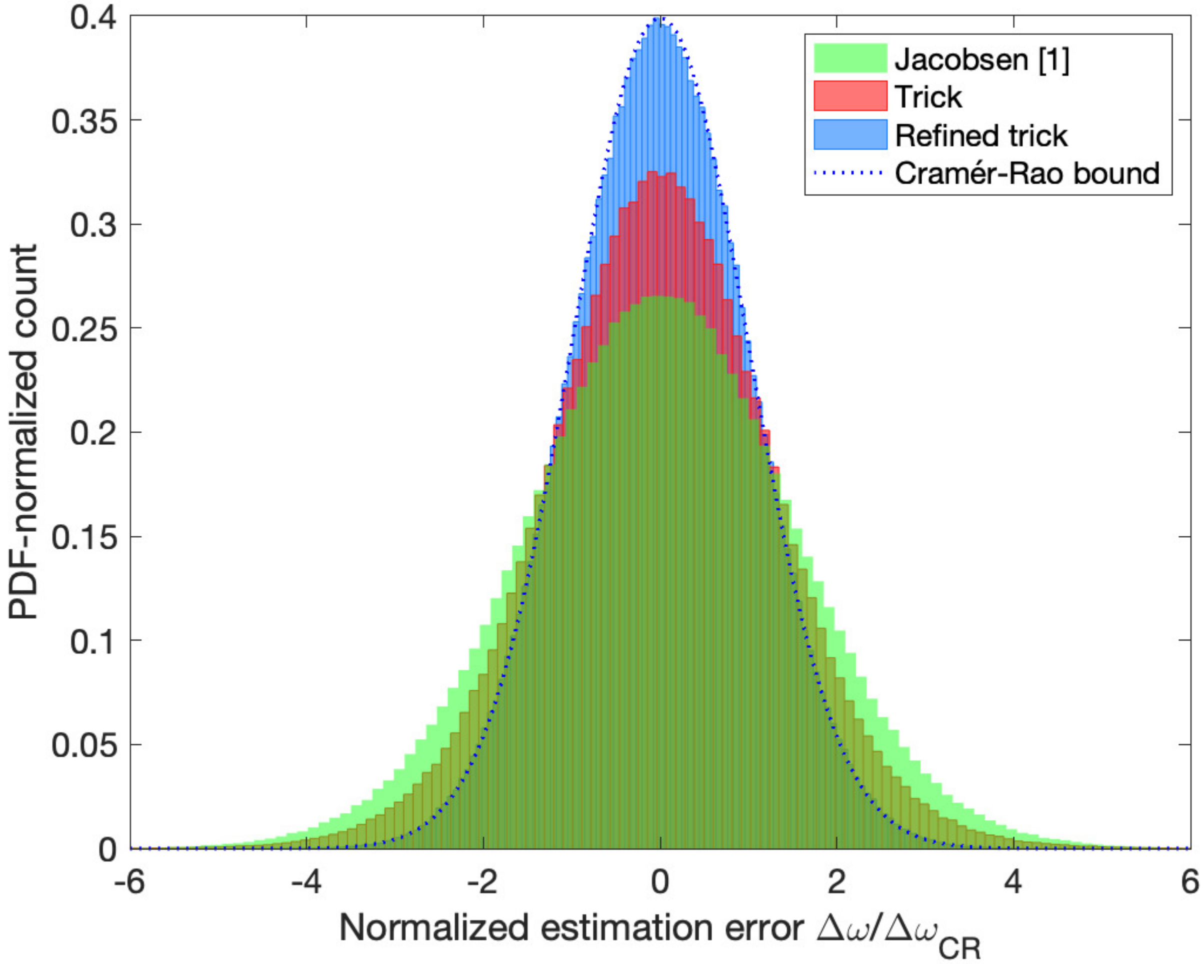

Fig. 3. Histograms from one million random tests (number of samples N, SNR, frequency, noise): the error that results from using Jacobsen's frequency estimator [1], the trick (1), or the refined trick (Fig. 2), is further normalized by the Cramér-Rao lower bound of the estimation problem ($\Delta\omega_{\text{CR}} = 2\sqrt{3}\,N^{-3/2}\text{SNR}^{-1}$). The standard deviations of the three estimators are 1.5325, 1.3008 and 1.0092, respectively.

An example of such calculation is shown in Appendix B, leading to the following statement: assume that the frequencies ω_k of the signal are distant from each other (modulo $(-\pi, \pi]$) by at least $\delta\omega > \pi/N$, and that the amplitude of the dominant sinusoid is A, then the estimation error of any of the frequencies of the signal is bounded according to

$$|\bar{\omega}_k - \omega_k| \leq \underbrace{\frac{2\pi}{N}}_{\text{DFT resolution}} \times \underbrace{\frac{2(K-1)\tan\big(\pi/(2N)\big)}{\pi\sin\big((\delta\omega - \pi/N)/2\big)}\frac{A}{|a_k|}}_{\text{“super-resolution” coefficient}}. \tag{4}$$

Here $\bar{\omega}_k, \omega_k, |a_k|$ are the estimated frequency, the ground-truth frequency and its amplitude, respectively. Despite its coarseness (see Fig. 4), this inequality already demonstrates super-resolution potential since the “super-resolution” coefficient is usually smaller than 1 and, in fact, tends to zero when N tends to infinity (for fixed $\delta\omega$).

Empirically, a minimum value of $4\pi/N$ for $\delta\omega$, or two DFT bins, seems to be sufficient to obtain good frequency estimates. Of course, this cheap approach to high-resolution multi-frequency estimation is not

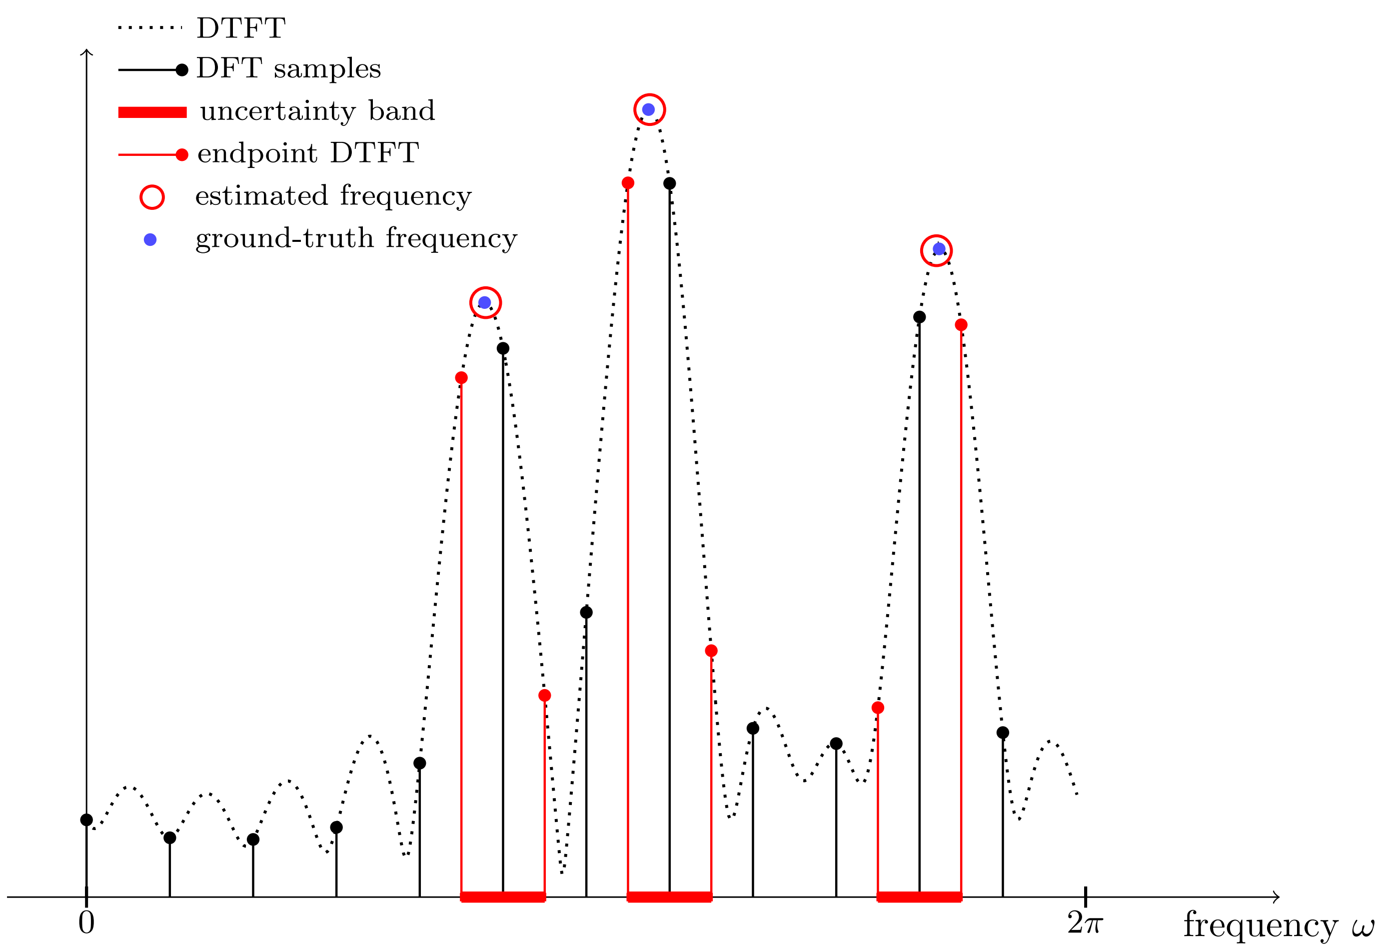

Fig. 4. Multiple frequencies (3 complex exponentials, 12 samples) can be accurately estimated by first locating the isolated peaks (e.g., using `findpeaks` in `Matlab`) and then applying (1) individually. The actual estimation errors of the three frequencies (from left to right) are roughly $(0.01, 0.02, 0.03) \times 2\pi/12$, well below the resolution $2\pi/12$ of the DFT; for comparison, the upper bound (4) provides the much more conservative values $(0.66, 0.45, 0.58) \times 2\pi/12$.

optimal; yet, it could be used as the starting point of any iterative algorithm designed to maximize the likelihood of the problem.

Conclusion

The frequency of a single complex exponential can be found exactly using the magnitude of only two samples of its DTFT, as this note shows. In the presence of noise or other inaccuracies, the trick that we provide is very robust, and can even be iterated once to reach the theoretical optimum (Cramér-Rao lower bound)—up to less than 1%. The robustness of this formula makes it possible to, e.g., refine the peaks of the DFT of a signal, but we also anticipate that it can be used as a tool for high-resolution frequency estimation. For teaching purposes, we provide a step-by-step proof which requires only undergraduate signal processing knowledge.

Box 2 Step-by-Step Didactic Proof of (1)

Steps

1) Geometric sum: $\sum_{n=0}^{N-1} z^n = \frac{z^N - 1}{z-1}$ with $z = \mathrm{e}^{j(\omega_0-\omega)}$ leads to $X(\omega) = a_0 \frac{\mathrm{e}^{jN(\omega_0-\omega)} - 1}{\mathrm{e}^{j(\omega_0-\omega)} - 1}$
2) Euler's formula: $\sin\theta = \frac{\mathrm{e}^{j\theta} - \mathrm{e}^{-j\theta}}{2j}$ leads to $\left| \begin{array}{l} X(\omega_0-\omega) = a_0 \mathrm{e}^{j\frac{N-1}{2}\omega} \frac{\sin\left(N\omega/2\right)}{\sin\left(\omega/2\right)} \\ \text{then } \left|X(\omega_0-\omega)\right| = |a_0| \left|\frac{\sin\left(N\omega/2\right)}{\sin\left(\omega/2\right)}\right| \end{array} \right.$
3) Notation: $\begin{cases} \omega_{12} = (\omega_1+\omega_2)/2 \\ B = (\omega_2-\omega_1)/4 \\ u = (\omega_0-\omega_{12})/2 \end{cases}$ leads to $\begin{cases} \left|X(\omega_1)\right| = |a_0| \left|\frac{\sin\left(N(u+B)\right)}{\sin(u+B)}\right| \\ \left|X(\omega_2)\right| = |a_0| \left|\frac{\sin\left(N(u-B)\right)}{\sin(u-B)}\right| \end{cases}$
4) π-periodicity of $|\sin x|$: $B = \text{integer} \times \frac{\pi}{2N}$ leads to $\frac{\left|X(\omega_2)\right|}{\left|X(\omega_1)\right|} = \left|\frac{\sin(u+B)}{\sin(u-B)}\right|$
5) Sign of sin: $\pm u \leq B$ leads to $\frac{\left|X(\omega_2)\right|}{\left|X(\omega_1)\right|} = \frac{\sin(B+u)}{\sin(B-u)}$
6) Trigonometry: $\sin(a \pm b) = \sin a \cos b \pm \cos a \sin b$ leads to $\frac{\left|X(\omega_2)\right|}{\left|X(\omega_1)\right|} = \frac{\tan B + \tan u}{\tan B - \tan u}$
7) Algebraic resolution: $\tan u = \tan B \times \frac{\left|X(\omega_2)\right| - \left|X(\omega_1)\right|}{\left|X(\omega_2)\right| + \left|X(\omega_1)\right|}$ which leads to

$$u = \frac{\omega_0 - \omega_{12}}{2} = \arctan\left(\tan\left(\frac{\omega_2-\omega_1}{4}\right) \frac{\left|X(\omega_2)\right| - \left|X(\omega_1)\right|}{\left|X(\omega_2)\right| + \left|X(\omega_1)\right|}\right)$$

Appendix — Other Proofs

A. Error bound (3)

A direct proof of this inequality uses the fact that $|\arctan a - \arctan b| \leq |a-b|$ and the triangle inequality $|a+b| \leq |a| + |b|$. More specifically, denoting by X_1, X_2 the DTFT of the "ground-truth" signal at ω_1, ω_2, and by ε_1, ε_2 the errors (caused by noise or otherwise) on $|X_1|$, $|X_2|$, we have

$$\begin{aligned} |\bar{\omega}_0 - \omega_0| &= \left| 2\arctan\left(\tan\left(\frac{\pi}{2N}\right) \frac{|X_1| + \varepsilon_1 - |X_2| - \varepsilon_2}{|X_1| + \varepsilon_1 + |X_2| + \varepsilon_2}\right) - 2\arctan\left(\tan\left(\frac{\pi}{2N}\right)\frac{|X_1| - |X_2|}{|X_1| + |X_2|}\right)\right| \\ &\leq 2\tan\left(\frac{\pi}{2N}\right) \left| \underbrace{\frac{|X_1| + \varepsilon_1 - |X_2| - \varepsilon_2}{|X_1| + \varepsilon_1 + |X_2| + \varepsilon_2} - \frac{|X_1| - |X_2|}{|X_1| + |X_2|}}_{= \frac{2\varepsilon_1\left(|X_2| + \varepsilon_2\right) - 2\varepsilon_2\left(|X_1| + \varepsilon_1\right)}{\left(|X_1| + |X_2|\right)\left(|X_1| + \varepsilon_1 + |X_2| + \varepsilon_2\right)}} \right| \\ &\leq \underbrace{2\tan\left(\frac{\pi}{2N}\right) \frac{\max\left(|2\varepsilon_1|, |2\varepsilon_2|\right)\left(|X_1| + \varepsilon_1 + |X_2| + \varepsilon_2\right)}{\left(|X_1| + |X_2|\right)\left(|X_1| + \varepsilon_1 + |X_2| + \varepsilon_2\right)}}_{= 4\tan\left(\frac{\pi}{2N}\right)\frac{\max\left(|\varepsilon_1|, |\varepsilon_2|\right)}{|X_1| + |X_2|}} \end{aligned}$$

which leads to the inequality (3) after noticing that $|X_1| + |X_2| \geq N|a_0|$ (because $\omega_0 \in [\omega_1, \omega_2]$).

B. Error bound (4)

Denoting by ω_1, ω_2,…ω_K the K different frequencies and a_1, a_2,…a_K the associated (complex-valued) amplitudes, the DTFT of the samples x_n is given by

$$X(\omega) = \sum_{k=1}^{K} a_k \frac{\mathrm{e}^{jN(\omega_k-\omega)}-1}{\mathrm{e}^{j(\omega_k-\omega)}-1}.$$

Evaluating the estimation error of the frequency ω_{k_0} using (1) requires calculating the bound ε in (3); i.e., the maximum error between $X(\omega)$ and the DTFT of a single frequency model, when $|\omega-\omega_{k_0}| \leq \pi/N$ (with the hypothesis that the minimum distance between ω_{k_0} and the other ω_k is at least $\delta\omega > \pi/N$):

$$\begin{aligned}
\left|X(\omega) - a_{k_0}\frac{\mathrm{e}^{jN(\omega_{k_0}-\omega)}-1}{\mathrm{e}^{j(\omega_{k_0}-\omega)}-1}\right| &= \left|\sum_{k\neq k_0} a_k \frac{\mathrm{e}^{jN(\omega_k-\omega)}-1}{\mathrm{e}^{j(\omega_k-\omega)}-1}\right| \\
&\leq \sum_{k\neq k_0} |a_k| \left|\frac{\sin\big(N(\omega_k-\omega)/2\big)}{\sin\big((\omega_k-\omega)/2\big)}\right| && \text{(using triangle inequality and Euler's formula)} \\
&\leq \sum_{k\neq k_0} \frac{A}{\left|\sin\big((\omega_k-\omega)/2\big)\right|} && \text{(denoting } A = \max_{k=1\ldots K}|a_k|\text{)} \\
&\leq \frac{(K-1)A}{\min_{k\neq k_0}\left|\sin\big((\omega_k-\omega)/2\big)\right|} \\
&\leq \frac{(K-1)A}{\sin\big((\delta\omega-\pi/N)/2\big)} && \text{(since } |\omega-\omega_{k_0}| \leq \pi/N < \delta\omega\text{).}
\end{aligned}$$

The right-hand side of the last inequality provides an upper bound for ε, that we can use in (3) to find

$$|\bar{\omega}_{k_0} - \omega_{k_0}| \leq \frac{2\pi}{N} \times \frac{2(K-1)\tan\big(\pi/(2N)\big)}{\pi\sin\big((\delta\omega-\pi/N)/2\big)} \frac{A}{|a_{k_0}|}.$$

References

[1] E. Jacobsen and P. Kootsookos, "Fast, accurate frequency estimators [DSP Tips & Tricks]," *IEEE Signal Processing Magazine*, vol. 24, no. 3, pp. 123–125, 2007.

[2] Ç. Candan, "A method for fine resolution frequency estimation from three DFT samples," *IEEE Signal processing letters*, vol. 18, no. 6, pp. 351–354, 2011.

[3] P. Stoica and A. Nehorai, "MUSIC, maximum likelihood, and Cramér-Rao bound," *IEEE Transactions on Acoustics, speech, and signal processing*, vol. 37, no. 5, pp. 720–741, 1989.

AUTHORS' BIOGRAPHIES

Ruiming Guo (ruiming.guo@imperial.ac.uk) received a B.E. degree in Electronic Engineering from Sichuan University, Chengdu, China, in 2017. He received his Ph.D. degree in Electronic Engineering at the Chinese University of Hong Kong (CUHK), HKSAR, China, in 2021. He worked as a postdoctoral research fellow with Prof. Thierry Blu at EE Department at CUHK, from 2021 to 2022. He is currently a postdoctoral research associate at the EEE Department of Imperial College London (ICL), working with Prof. Ayush Bhandari on computational imaging and modulo sampling. He received the Postgraduate Student Research Excellence Awards from the Department of the Electronic Engineering of the Chinese University of Hong Kong in 2022. His research interests include sparse signal processing, sampling theory, inverse problems, modulo sampling, and computational sensing and imaging.

Thierry Blu (thierry.blu@m4x.org) received the "Diplôme d'ingénieur" from École Polytechnique, France, in 1986 and the Ph.D from Télécom Paris (ENST), in 1996. He has been with the Department of Electronic Engineering, the Chinese University of Hong Kong, since 2008. He received two best paper awards (2003 and 2006), and is coauthor of a paper which received a Young Author best paper award (2009), all from the IEEE Signal Processing Society. He is a Fellow of the IEEE (2012). His research interests include wavelets, approximation and sampling theory, sparse representations, biomedical imaging, optics, and wave propagation.