

A correspondence between zeros of time-frequency transforms and Gaussian analytic functions

Rémi Bardenet, Pierre Chainais
Université de Lille, CNRS,
Centrale Lille, UMR 9189 – CRISTAL,
59651 Villeneuve d’Ascq Cedex, France.
Email: remi.bardenet@gmail.com

Julien Flamant
Université de Lorraine
CNRS, CRAN
54000 Nancy, France

Adrien Hardy
Université de Lille,
Laboratoire Paul Painlevé, UMR 8524,
59655 Villeneuve d’Ascq Cedex, France.

Abstract—In this paper, we survey our joint work on the point processes formed by the zeros of time-frequency transforms of Gaussian white noises [1], [2]. Unlike both references, we present the work from the bottom up, stating results in the order they came to us and commenting what we were trying to achieve. The route to our more general results in [2] was a sort of ping pong game between signal processing, harmonic analysis, and probability. We hope that narrating this game gives additional insight into the more technical aspects of the two references. We conclude with a number of open problems that we believe are relevant to the SampTA community.

Index Terms—time-frequency/time-scale analysis, white noise, spectrogram zeros, point process

I. INTRODUCTION

Spectrograms are a cornerstone of time-frequency analysis [3]. They are quadratic time-frequency representations of a signal [4, Chapter 4], associating to each time and frequency a real number that measures the energy content of a signal at that time and frequency, unlike global-in-time tools such as the Fourier transform. Since it is natural to expect that there is more energy where there is more information or signal, most methodologies have focused on detecting and processing the local maxima of the spectrogram; see e.g. [3]. Usual techniques include *ridge extraction*, e.g., to identify chirps, *reassignment* and *synchrosqueezing*, to better localize the maxima before further quantitative analysis.

In contrast, Patrick Flandrin has recently observed [5] that the locations of the zeros of a spectrogram in the time-frequency plane almost completely characterize the spectrogram, and he proposed to use the point pattern formed by the zeros in filtering and reconstruction of signals in noise. This proposition stems from the empirical observation that the zeros of the short-time Fourier transform (STFT) of white noise are uniformly spread over the time-frequency plane, and tend not to cluster, as if they repelled each other. In the presence of a signal, zeros are absent in the time-frequency support of the signal, thus creating large holes that appear to be very rare when observing pure white noise. This leads to testing the presence of signal by looking at statistics of the point pattern of zeros, and trying to identify holes.

When we first heard of [5], we were working on Monte Carlo integration with repulsive point processes [6]. More precisely, we were working on orthogonal polynomial ensembles, a family of determinantal point processes (DPPs). Loosely speaking, a point process is a random discrete set of points. DPPs were introduced by Odile Macchi in 1975 [7] as models for physical beams of fermions. They are particularly tractable repulsive point processes parametrized by a kernel, and their analysis crucially involves *reproducing* kernels [8]. The similarity of this mathematical backbone of DPPs with the analytical tools behind time-frequency analysis [4] made us conjecture that the point process formed by the zeros of the STFT of white Gaussian noise was a DPP. As we explain in Section II, this is not the case. However, once all the maths were laid out to characterize the zeros of random spectrograms, we realized that we could have obtained a well-known DPP if we had considered the analytic wavelet transform of [9] instead of the STFT; this is Section III. Then again, we recognized the key role implicitly played by generating functions of orthogonal polynomials. This allowed us to further generalize our results on the identification of zeros to more time-frequency transforms; see Section IV. Finally, we discuss open questions for the SampTA audience in Section V.

II. FROM THE STFT TO THE PLANAR GAF

Let $f \in L^2(\mathbb{R})$ and g be the pdf of a centered Gaussian, with variance normalized so that $\|g\|_2 = 1$. The evaluation at $(u, v) \in \mathbb{R}^2$ of the short-time Fourier transform (STFT) of f with Gaussian window g reads

$$V_g f(u, v) = \int f(t) \overline{g(t-u)} e^{-2i\pi tv} dt = \langle f, M_v T_u g \rangle, \quad (1)$$

with $\langle \cdot, \cdot \rangle$ denoting the inner product in $L^2(\mathbb{R})$, $M_v f = e^{2i\pi v \cdot} f(\cdot)$ and $T_u f = f(\cdot - u)$.

Our first step was to rigorously define the STFT of complex white Gaussian noise (WGN), in a manner that both satisfied the intuition we may have of white noise, and allowed us to speak of the zeros of the spectrogram. In [1, Section 3.1], we used a classical definition of real WGN as a random tempered distribution with Gaussian characteristic function. A complex WGN ξ is then defined as having independent white Gaussian noises ξ_1, ξ_2 as its real and imaginary parts.

This allowed us to *pointwise* compute $\langle \xi, M_v T_u g \rangle$ since $M_v T_u g$ is smooth. By pointwise, we mean for a given time u and frequency v . Then in [1, Proposition 3], we stitched all these pointwise evaluations together and concluded that with probability 1 on the white noise, the zeros of the function

$$z = u + iv \mapsto \langle \xi, M_v T_u g \rangle$$

are those of the entire function

$$\text{GAF}_{\mathbb{C}}(z) := \sum_{k=0}^{\infty} a_k \frac{\pi^{k/2} z^k}{\sqrt{k!}}, \quad (2)$$

where $a_k = (\langle \xi_1, h_k \rangle + i \langle \xi_2, h_k \rangle) / \sqrt{2}$ are i.i.d. unit complex Gaussians, and h_k are the Hermite functions. To read the current paper, it is enough to know that Hermite functions are a basis of $L^2(\mathbb{R})$ built using orthogonal polynomials with respect to the Gaussian g .

All the above was the formalisation of the following heuristic. A complex WGN should have unit complex Gaussian coefficients in any basis of $L^2(\mathbb{R})$. Among such bases, we pick the Hermite basis, because Hermite functions essentially map to complex monomials through the STFT [4]. We thus expect an analytic function of the form (2) as the STFT of ξ .

Identifying the zeros of the STFT of white noise with the zeros of a random entire function on the time-frequency plane is satisfactory, since entire functions have isolated zeros. It thus makes sense to speak of the point process of the zeros. Furthermore, the zeros of (2) are actually a well-studied stationary point process [10], called the zeros of the planar Gaussian analytic function (GAF), of which we show a sample in Figure 1. Combining known probabilistic results on this point process and standard statistical methodology for spatial point processes, we obtained in [1, Section 5] powerful signal detection procedures that strengthen the seminal work of [5].




Fig. 1. White dots are a sample zero set of the planar GAF, superimposed on a Charlier spectrogram, see [2] for the definition of the Charlier transform.

To conclude, one can show [10, Section 5.1] that the zeros of the planar GAF are not a DPP with Hermitian kernel. Strictly speaking, they could still be a DPP with a nonhermitian kernel.

III. FROM A DPP TO ANALYTIC WAVELETS

We were slightly frustrated of not finding a well-known DPP behind the STFT of white noise in [1]. In particular, for $\alpha > -1$, consider the random analytic function

$$\text{GAF}_{\mathbb{D}}^{(\alpha)}(z) := \sum_{k \in \mathbb{N}} a_k \sqrt{\frac{\Gamma(k + \alpha + 1)}{k!}} z^k, \quad |z| < 1, \quad (3)$$

where the a_i are again i.i.d. unit complex Gaussians. The series (3) defines the so-called *hyperbolic GAF* [10]. When $\alpha = 0$, it simply reads $\sum a_k z^k$ and its zeros are a DPP. A natural¹ question is then: what signal processing problem should we have tried to solve in order to end up with the zeros of the hyperbolic GAF instead of the planar GAF? The gamma coefficients made us think of replacing Hermite functions by Laguerre functions, and insight from Patrick Flandrin helped us connect to the analytic wavelet transform. As we shall now see, this backward reasoning turned out to be mathematically valid and was the starting point of our next work [2].

The same heuristic as in Section II holds: we need a basis of some Hilbert space of signals, a white noise that can be decomposed onto that basis, and a time-frequency transform that sends elements of the basis onto the right monomials.

A. The space of signals

Laguerre functions are a basis of $L^2(\mathbb{R}_+)$ built using the orthogonal polynomials with respect to the pdf of a gamma random variable. Thus, we are looking for a space of signals where the half-line \mathbb{R}_+ plays a role. A natural² example is the space of “analytic” signals

$$H^2(\mathbb{R}) := \{f \in L^2(\mathbb{R}, \mathbb{C}) : \text{Supp}(\hat{f}) \subset \mathbb{R}_+\}.$$

The inverse Fourier transforms (f_k) of the Laguerre functions form a basis of $H^2(\mathbb{R})$.

B. The time-scale transform

Now, we need a transform that maps the basis (f_k) to monomials. Let $\beta > -1/2$ and set

$$\psi_\beta(t) := \frac{1}{(t + i)^\beta}, \quad t \in \mathbb{R}.$$

The Fourier transform of ψ_β is essentially a gamma pdf

$$\frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}_+} x^\beta e^{-x} e^{itx} dx = -\frac{1}{\sqrt{2\pi}} \frac{(-i)^{\beta+1}}{(t + i)^{\beta+1}} \Gamma(\beta + 1),$$

so that ψ_β is intuitively a fine window if we want to use Laguerre calculus in the Fourier domain. The Daubechies-Paul wavelet transform of $f \in H^2(\mathbb{R})$ is defined as

$$W_\beta f(u, s) := \langle f, T_u D_s \psi_\beta \rangle, \quad (4)$$

where $u \in \mathbb{R}$ and $s \in \mathbb{R}_+^*$ are thought of as time and scale, $T_u f(t) := f(t - u)$ and $D_s f(t) := s^{-1/2} f(t/s)$ are the usual translation and dilation operators.

¹At least to DPP aficionados.

²this time, to signal processing fans!

We show in [2, Theorem 2.3] that for $f \in H^2(\mathbb{R})$ and up to a nonvanishing term,

$$W_\beta f(-u, s) \propto \mathcal{L}_\mathbb{D}^{(\beta)} f(\varphi(u + is)), \quad (5)$$

where $u + is \in \mathbb{C}_+$ is in the time-scale half-plane, $\varphi(w) := (2w - i)/(2w + i)$ is a conformal mapping from \mathbb{C}_+ to \mathbb{D} , and $\mathcal{L}_\mathbb{D}^{(\beta)}$ is a transform that precisely sends the inverse Fourier transforms of the Laguerre functions onto the right monomials

$$\mathcal{L}_\mathbb{D}^{(\beta)} f_k(z) = \sqrt{\frac{\Gamma(k + 2\beta + 1)}{k!}} z^k. \quad (6)$$

Everything is in place but the white noise. Heuristically, if we were able to somehow decompose a white noise ξ on $H^2(\mathbb{R}_+)$ as $\sum_{k \geq 0} a_k f_k$, where a_k are i.i.d. unit complex Gaussians, then (5) and (6) should help us conclude that

$$W_\beta \xi(z) \propto \sum_{k \geq 0} a_k \sqrt{\frac{\Gamma(k + 2\beta + 1)}{k!}} z^k = \text{GAF}_\mathbb{D}^{(2\beta)}(z)$$

up to nonvanishing terms. This would characterize the zeros of the analytic wavelet transform of white noise.

C. The white noise

The series $\sum_{k \geq 0} a_k f_k$ almost surely diverges in $H^2(\mathbb{R})$. To remedy this, we follow [11] in his formalization of abstract Wiener spaces. First, we define a slightly weaker norm on $H^2(\mathbb{R})$ that ensures the series converges

$$\|f\|_\Theta^2 := \sum_{k \in \mathbb{N}} \frac{1}{1 + k^2} |\langle f, f_k \rangle|^2. \quad (7)$$

Indeed, it follows that

$$\left\| \sum_{k \in \mathbb{N}} a_k f_k \right\|_\Theta^2 = \sum_{k \in \mathbb{N}} \frac{|\xi_k|^2}{1 + k^2} < \infty \quad \text{a.s.} \quad (8)$$

since the right hand side has finite expectation. Now define Θ as the completion of $H^2(\mathbb{R})$ for the norm (7). Leonard Gross' construction [11] then gives a unique probability measure on Θ , called white Gaussian noise, such that

$$\int_\Theta e^{i\lambda(\theta)} \mu(d\theta) = e^{-\|h_\lambda\|^2/2}, \quad \lambda \in \Theta^*, \quad (9)$$

where Θ^* is the topological dual of Θ and $h_\lambda \in H^2(\mathbb{R})$ is associated to $\lambda \in \Theta^* \subset H^2(\mathbb{R})$ by Riesz' representation theorem. Equation (9) is to be interpreted as specifying the Fourier transform of the white noise on $H^2(\mathbb{R})$. In particular, it leads to $\sum a_k f_k$ converging almost surely in Θ to μ , which is our goal. We can even give rates of convergence [2, Section 5]. We also note in passing that μ is a rigorous definition of an analytic white noise.

With the white noise now defined, we prove in [2, Theorem 2.3] that $\mathcal{L}_\mathbb{D}^{(\beta)} \xi$ can be unambiguously defined as a random analytic function in the unit open disk \mathbb{D} . Furthermore, it is equal in distribution to $\text{GAF}_\mathbb{D}^{(2\beta)}$. In particular, the case $\beta = 0$ yields zeros that are a DPP! A sample of this DPP obtained as the numerical zeros of a Paul-Daubechies scalogram is shown in Figure 2.

Finally, we note that Abreu, Haimi, Koliander, and Romero [12] have also independently proved that the zeros of $W_\beta \xi$ are the zeros of the hyperbolic GAF, with a slightly different definition of white noise.

IV. THE GENERAL PATTERN

Having extended the STFT case to the analytic wavelet transform, we then realized that the mathematical construction in Section III was much more general than these two cases. The whole idea is to first fix a Hilbert space \mathcal{H} of signals, fix an orthonormal basis (f_k) of \mathcal{H} , and complete \mathcal{H} into Θ as in Section III-C, so as to be able to define a white Gaussian noise on Θ . Then pick a sequence $(\Psi_k)_{k \in \mathbb{N}}$ of holomorphic functions on a domain Λ satisfying

$$\sup_{z \in K} \sum_{k \in \mathbb{N}} |\Psi_k(z)|^2 < \infty \quad (10)$$

for any compact $K \subset \Lambda$, and define \mathcal{L} acting on \mathcal{H} as

$$\mathcal{L}f(z) := \sum_{k=0}^{\infty} \langle f_k, f \rangle \Psi_k(z), \quad z \in \Lambda. \quad (11)$$

In signal processing terms, the signal $f \in \mathcal{H}$ is analyzed by \mathcal{L} in the basis (f_k) , and reconstructed as an analytic function on Λ using the dictionary Ψ_k . Note that dominated convergence allows us to rewrite \mathcal{L} in kernel form,

$$\mathcal{L}f(z) := \langle T(\cdot, z), f \rangle, \quad T(x, z) := \sum_{k=0}^{\infty} f_k(x) \Psi_k(z). \quad (12)$$

The construction of white noise in Section III-C guarantees that \mathcal{L} extends uniquely to Θ and that $\mathcal{L}\xi$ has the same law as the GAF $\sum a_k \Psi_k$; see our generic [2, Theorem 3.4].

The STFT with Gaussian window and the analytic wavelet transform of [9] are particular cases of \mathcal{L} , where the representation (12) is provided by identities for the generating functions T of classical orthogonal polynomials. This prompted us to plug other generating functions for orthogonal polynomials as kernels, and check whether the resulting transform had a time-frequency/scale interpretation.

Discrete orthogonal polynomials, for instance, lead to *discrete* transforms, that result in a transform of WGN that is either the planar (2), the hyperbolic (3), or the *spherical* GAF

$$\text{GAF}_\mathbb{S}^{(N)}(z) := \sum_{k=0}^N \xi_k \sqrt{\binom{N}{k}} z^k, \quad z \in \mathbb{C}. \quad (13)$$

We plot in Figure 3 a sample of the spherical GAF and its zeros, mapped onto a sphere through stereographic projection. This rendering highlights the invariance of the set of zeros to isometries of the sphere, just like the zeros of the planar and hyperbolic GAFs are invariant to isometries of the plane and the Poincaré half-plane, respectively.

These discrete time-frequency transforms are *not* discrete approximations of either the STFT or the analytic wavelet transforms, yet they preserve the distribution of the zeros of the corresponding spectrograms. If zeros end up competing with maxima in signal processing practice, these discrete transforms




Fig. 2. White dots are a sample zero set of the hyperbolic GAF, superimposed on an analytic scalogram of an analytic white noise.




Fig. 3. White dots are a sample zero set of the spherical GAF, superimposed on a Krawtchouk spectrogram, see [2] for definitions.

may be of computational interest. Indeed, they remove the need to discretize continuous Fourier transforms.

V. OPEN ISSUES

On the unicity of GAFs. There are some fundamental limitations to the correspondence between GAFs and time-frequency transforms. The three GAFs we introduced – planar, hyperbolic, and spherical – are canonical in the sense that they are the only GAFs that are invariant to isometries of the complex plane, the Poincaré half-plane, and the sphere. From a statistical point of view, such invariances are crucial, as they make the point process of the zeros stationary and allow estimating function statistics [1, Section 5]. Extending the correspondence will require either more exotic geometries, or dropping the Gaussian/analytic aspects on the GAF side.

On the unicity of TF transforms. Another point in favour of removing analyticity is that the Gaussian window is the only window that makes the STFT map $L^2(\mathbb{R})$ to analytic functions [13]. The analytic wavelet transform of Section III plays the same unique role among continuous wavelet transforms [14].

Nevertheless, our discrete transforms in [2] do bypass this limitation while still mapping to analytic functions.

Beyond orthogonal polynomials. The transforms that we discussed in this paper are related to orthogonal polynomials. However, we could easily generalize to other bases of L^2 , provided that we can associate them with a dictionary (Ψ_k) of holomorphic functions such that the kernel (12) makes sense for time-frequency/scale analysis.

On maxima. We briefly discuss local extrema of the STFT of white noise in [2, Corollary 2.8], and we show that they are the zeros of a GAF, although this is not the planar GAF, so that local extrema do not form a stationary point process. It seems as if zeros have more structure, which is an argument in favour of using them in filtering as suggested by [5], but a more thorough study of maxima is needed.

REFERENCES

- [1] R. Bardenet, J. Flamant, and P. Chainais, “On the zeros of the spectrogram of white noise,” *ACHA*, 2018.
- [2] R. Bardenet and A. Hardy, “Time-frequency transforms of white noises and Gaussian analytic functions,” *Arxiv preprint:1807.11554*, 2018.
- [3] P. Flandrin, *Time-frequency/time-scale analysis*. Academic press, 1998.
- [4] K. Gröchenig, *Foundations of time-frequency analysis*. Birkhäuser, 2001.
- [5] P. Flandrin, “Time–frequency filtering based on spectrogram zeros,” *IEEE Signal Processing Letters*, vol. 22, no. 11, pp. 2137–2141, 2015.
- [6] R. Bardenet and A. Hardy, “Monte Carlo with determinantal point processes,” *arXiv preprint arXiv:1605.00361*, 2016.
- [7] O. Macchi, “The coincidence approach to stochastic point processes,” *Advances in Applied Probability*, vol. 7, pp. 83–122, 1975.
- [8] J. B. Hough, M. Krishnapur, Y. Peres, and B. Virág, “Determinantal processes and independence,” *Probability surveys*, 2006.
- [9] I. Daubechies and T. Paul, “Time-frequency localization operators – a geometric phase space approach: II. the use of dilations,” *Inverse problems*, vol. 4, pp. 661–680, 1988.
- [10] J. B. Hough, M. Krishnapur, Y. Peres, and B. Virág, *Zeros of Gaussian analytic functions and determinantal point processes*. American Mathematical Society, 2009, vol. 51.
- [11] L. Gross, *Abstract Wiener Spaces*. Proc. Fifth Berkeley Sympos. Math. Statist. and Probability (1965/66), Vol. II. Univ. California Press., 1967.
- [12] L. D. Abreu, A. Haimi, G. Koliander, and J. L. Romero, “Filtering with wavelet zeros and Gaussian analytic functions,” *arXiv preprint arXiv:1807.03183*, no. v1, 2018.
- [13] G. Ascensi and J. Bruna, “Model space results for the Gabor and Wavelet transforms,” *IEEE Transactions on Information Theory*, 2009.
- [14] N. Holighaus, G. Koliander, Z. Průša, and L. D. Abreu, “Characterization of analytic wavelet transforms and a new phaseless reconstruction algorithm,” *To appear in IEEE Transactions on Signal processing*, 2019.