

THE ODD FIBONACCI ZETA FUNCTION

DAVID LOWRY-DUDA

ABSTRACT. We look at the (odd) Fibonacci zeta function and comment briefly on the behavior of its zeros and poles.

1. INTRODUCTION

In forthcoming work [AKLDW24a, AKLDW24b] with my collaborators Eran, Chan, and Alex, we describe different ways to understand the Fibonacci zeta function

$$Z_{\text{Fib}}(s) := \sum_{n \geq 1} \frac{1}{F(n)^s} = 1 + 1 + \frac{1}{2^s} + \frac{1}{3^s} + \dots$$

and certain generalizations associated to real quadratic fields. I gave a talk [LD20b] about this in 2020 and it's sat in various states since then.

The focus of our forthcoming papers is that these objects (and their generalizations) can be understood in several different ways — including one way through modular forms! This is surprising and interesting.

This note is not as surprising nor as deep. In this focused technical note, I consider the question of studying the zeros and plot of the Fibonacci zeta function.

Actually, I restrict my attention to the *odd Fibonacci zeta function*

$$Z(s) := Z_{\text{Fib}}^{\text{odd}}(s) = \sum_{n \geq 1} \frac{1}{F(2n-1)^s},$$

consisting of odd-indexed terms of the Fibonacci zeta function. Although it's possible to study the whole Fibonacci zeta function (or alternately just the even-indexed terms, the *even Fibonacci zeta function*), the odd Fibonacci zeta function has a simple analytic description and behaves in many ways like a simplest component.

In our work, we prove the following continuation.

Theorem 1. *Write $\varepsilon = \frac{1+\sqrt{5}}{2}$. Then for all $s \in \mathbb{C}$ away from poles of the summands,*

$$Z(s) := Z_{\text{Fib}}^{\text{odd}}(s) = \frac{5^{s/2}}{8\Gamma(s)\log \varepsilon} \sum_{m \in \mathbb{Z}} (-1)^m \Gamma\left(\frac{s}{2} + \frac{\pi i m}{2\log \varepsilon}\right) \Gamma\left(\frac{s}{2} - \frac{\pi i m}{2\log \varepsilon}\right). \quad (1)$$

I take this for granted and defer the proof to our forthcoming work.

This continuation shows many properties. There are trivial zeros (coming from poles of $\Gamma(s)$) that are reminiscent of the trivial zeros of the standard Riemann zeta function, $\zeta(s)$; and there is a half-lattice of poles from the other Gamma functions. Other than that, this is a Dirichlet series with meromorphic continuation.

Date: December 4, 2024.

This work was supported by the Simons Collaboration in Arithmetic Geometry, Number Theory, and Computation via the Simons Foundation grant 546235.

It is easy to write down Dirichlet series, but a random Dirichlet series will almost surely have a natural boundary on the abscissa of convergence. Dirichlet series with continuation are special.

On the other hand, I've almost always studied *automorphic L*-functions, which are the nicest possible family of Dirichlet series.

We examine two guiding questions in this note:

- (1) What does $Z(s)$ look like?
- (2) Where are its zeros?

Through this exploration, we might build up additional intuition. If we only look at the nicest possible Dirichlet series, we can forget what makes them special or what properties actually distinguish them.

2. MAKING PLOTS

From (1) and Stirling's formula, it's clear that the summands decay extremely quickly in m . Using only the first $2|\text{Im } t|$ summands will be sufficient. I always use at least 20 to prevent over approximating. Thus making a plot of $Z(s)$ reduces to quickly computing lots of Gamma functions, adding them together, and then visualizing the result.

I typically use Sagemath [Sag20] for plots of complex functions (especially as I wrote the complex plotting routines in [LD20a]). But it turns out that sage is *stunningly slow* at making nice plots of $Z(s)$.

So instead, I describe here a simple implementation in C++ (with plotting in python). This takes two steps.

- (1) First, compute an approximation of $Z(s)$ on a square grid of points.
- (2) Then import this grid of points into python for plotting.

2.1. Computing $Z(s)$ on a mesh. The only nontrivial computational component is computing the gamma function at complex arguments. There are many implementations, but I tend to use `arblib` (now part of `flintarb`) [Joh17] for all sorts of special functions, unless it's not sufficient for some reason.

Remark 2. `arb` gives provably strong results and typically works in interval arithmetic. This takes additional work. Here I immediately throw away the intervals and just use the midpoint approximations. This is wasteful, but `arb` is fast and correct. What more could I ask for?

I use `acb.h` and `acb_gamma` from `arb`.

```

1 #include <complex>
2 #include <acb.h>
3 #include <arb.h>
4 #include <arf.h>
5
6 std::complex<double> gamma(const std::complex<double>& z) {
7     const slong prec = 53;
8     acb_t zz, res; acb_init(zz); acb_init(res);
9     acb_set_d_d(zz, z.real(), z.imag());
10    acb_gamma(res, zz, prec);
11    double real = arf_get_d(
12        arb_midref(acb_realref(res)), ARF_RND_NEAR
13    );
14    double imag = arf_get_d(
15        arb_midref(acb_imagref(res)), ARF_RND_NEAR

```

```

16 );
17 acb_clear(zz); acb_clear(res);
18 return std::complex<double>(real, imag);
19 }

```

In practice, `gamma` shadows a common `map`, so I actually put this in my own namespace `dld::gamma`.

Computing $Z(s)$ is now completely routine. Take m up to $\max(20, \lceil |\text{Im}(s)| \rceil)$ in (1) and add them up. Complete code is included in the appendix.

Making the Mesh. I use a slightly nontrivial way to compute the grid itself—nontrivial because I use multithreading to make this a parallel computation.

The overall structure is simple: divide the rows into number-of-threads many groups and then assign appropriate rows to each thread. For simplicity I have each thread write its rows to a separate temporary file and then combine them together afterwards.

This is my extremely simple form of multithreading. If some rows are far more computationally intense (which they are), then this won't evenly distribute computational load. But the naive independence means that I don't have to worry about shared memory or mutexes or other concurrency problems. A small possible improvement would be to split into more chunks and have a thread pool, but that added complexity doesn't seem worth it to me.

Regardless, the output is a CSV containing (x, y, value) rows. I use two variants: one where the values are $\arg(Z(x + iy))$ and the other where the values are the pair $(\text{Re}(Z(x + iy)), \text{Im}(Z(x + iy)))$

2.2. Plotting the data. Perhaps the most obvious way to plot the resulting data would be to assign your domain coloring (cf. [LD21], or https://en.wikipedia.org/wiki/Domain_coloring, or the beautiful visualizations of Frank Farris) and then color each computed point. Unfortunately, this requires an enormously dense grid of computed pixels to make a good looking plot.

Instead, we interpolate between computed pixels. And by “we”, I mean `matplotlib`, as this is implemented with various interpolation algorithms in `imshow`.

Alternately, to plot lines when the real or imaginary parts are 0, we can use `matplotlib`'s `contour`, which implements a marching squares algorithm for contour evaluation.

Both approaches can introduce artifacts. Fortunately, we are plotting holomorphic functions and the local behavior is either tame or too complex for any plotting method to have a chance (in practice). And in practice the viewer can detect when an artifact is introduced because it looks funny.

The broad look of these plotting routines in python look like the following.

```

1 # x_vals, y_vals, arg_vals, real_vals, imag_vals from CSV
2 import matplotlib.pyplot as plt
3
4 ## I stored these in most significant bit order, hence 'F'
5 arg_grid = arg_vals.reshape((grid_size, grid_size), order='F')
6 plt.imshow(arg_grid,
7             extent=[x_vals.min(), x_vals.max(),
8                     y_vals.min(), y_vals.max()],
9             origin='lower', interpolation="catmul")
10
11 # Or, for contours #
12 plt.contour(imag_grid, colors="#f9ae54",

```

```

13     extent=[x_vals.min(), x_vals.max(),
14             y_vals.min(), y_vals.max()],
15             origin="lower")
16 plt.contour(real_grid, colors="#0482d7",
17             extent=[x_vals.min(), x_vals.max(),
18                     y_vals.min(), y_vals.max()],
19                     origin="lower")

```

These produce the images in Figure 1. The argument plot in the Figure uses a colormap that is discontinuous at the boundary, which (in my normalization) amounts to a light-dark discontinuity when the imaginary part is 0.

3. COMMENTS ON THE PLOTS AND ZEROS

I chose to plot in the rectangle with opposite vertices $-35 - 5i$ and $5 + 35i$. The aspect ratio is 1 because I almost always want my aspect ratio to be 1. The odd Fibonacci zeta function $Z(s)$ is rather boring for $\operatorname{Re} s > 0$ and antisymmetric over the real axis, so it's not worth including too much in those directions.

We know that $Z(s)$ has a half-lattice of poles, so we expect the lattice like regularity. We also understand the alternating zeros and poles on the real line.

But it turns out that there seems to be approximately one zero rather close to every pole.

Initially I thought this was surprising. But the zero and polar behavior is largely constrained by Jensen's Formula. I state it in two forms (both of which I learned as Jensen's Formula as a grad student).

Theorem 3 (Jensen's Formula V1). *Let $f \neq 0$ be meromorphic on the closed disk $B_R(0)$. Let a_1, \dots, a_p denote the zeros of f in $B_R(0)$, counting multiplicities, and let b_1, \dots, b_q denote the poles in $B_R(0)$, also with multiplicities. Then for any z in $|z| < R$ which is not a zero or a pole,*

$$\begin{aligned} \log|f(z)| = & \int_0^{2\pi} \frac{R^2 - |z|^2}{|Re^{i\theta} - z|^2} \log|f(Re^{i\theta})| \frac{d\theta}{2\pi} \\ & - \sum_{i=1}^p \log \left| \frac{R^2 - \overline{a_i}z}{R(z - a_i)} \right| + \sum_{j=1}^q \log \left| \frac{R^2 - \overline{b_j}z}{R(z - b_j)} \right|. \end{aligned}$$

Theorem 4 (Jensen's Formula V2). *With the same notation as above, except omitting zeros or poles at $z = 0$; let $f(z) = c_f z^{\operatorname{ord}(0)} + \dots$, where c_f is the leading nonzero coefficient of the Laurent expansion around 0. Then*

$$\log|c_f| = \int_0^{2\pi} \log|f(Re^{i\theta})| \frac{d\theta}{2\pi} - \sum_{i=1}^p \log \left| \frac{R}{a_i} \right| + \sum_{j=1}^q \log \left| \frac{R}{b_j} \right| - (\operatorname{ord}(0)) \log R. \quad (2)$$

The second form comes essentially from expanding at $z = 0$ after modifying and removing the zero/polar behavior there.

The point is that $Z(s)$ grows too slowly to have interesting polar or zero behavior. Rearranging (2) and estimating growth, we observe that

$$- \sum_{i=1}^p \log \left| \frac{R}{a_i} \right| + \sum_{j=1}^q \log \left| \frac{R}{b_j} \right| \ll \log R.$$

(I use \ll here to mean that the LHS is bounded in magnitude and not merely bounded above). There are approximately R^2 poles within R of the origin, in a half-lattice. In order to counteract this growth, there must be regularly spaced zeros as well.

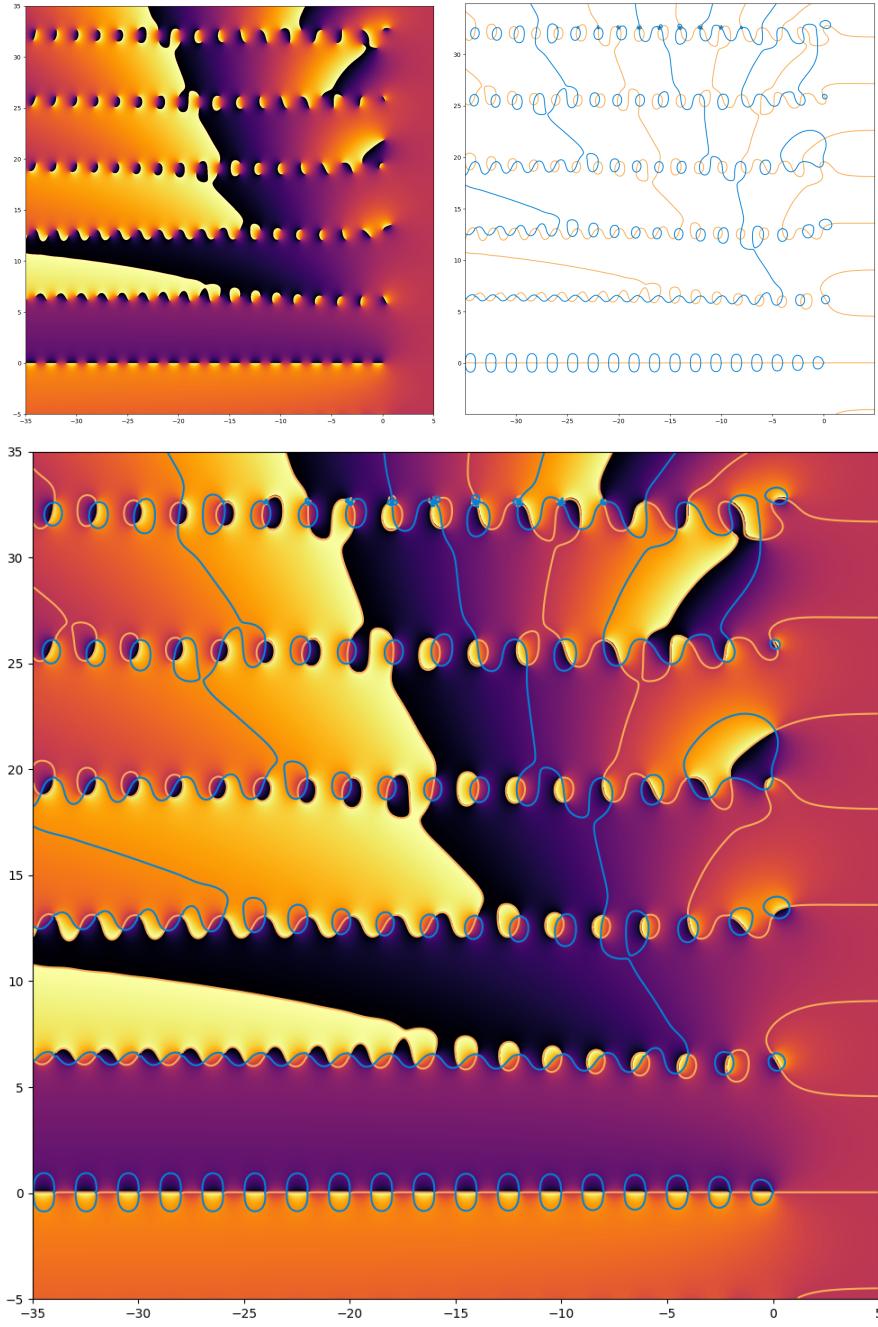


FIGURE 1. Interpolated imshow and contour plots of $Z(s)$. The last plot is the combined plot.

It's probably possible to say more about the distribution. I suspect that performing argument-principle type evaluation would show that the difference in the number of zeros and poles in a circle is bounded by $O(\log R)$, but I don't carry out this analysis.

3.1. Specific Zeros. Do the zeros have any meaning? Should we be able to predict them individually, as opposed to in distribution?

As a first step, we compute a couple of zeros using Newton's method in Figure 2.

0.41054841	+	5.80204742i
0.82209560	+	13.5268604i
-0.32303194	+	19.0607822i
0.13195381	+	21.8573826i
0.23103157	+	25.7125161i

FIGURE 2. The first five zeros corresponding to the first five poles of $Z(s)$.

We make several remarks.

First, these seem to be where we expect from the plots.

Second, there aren't obvious patterns or regularities in the real or imaginary parts (aside from being spaced about as far apart as the poles). Conceivably, these could be algebraic over $\mathbb{Q}[\sqrt{5}, \log(\varphi), \pi]$. Small tests don't support this. (It would not be hard to compute these and many more to much higher precision, and then do a more sophisticated analysis. I don't do this).

Third, several zeros occur in the region of absolute convergence of the underlying Dirichlet series. As Dirichlet series in the region of absolute convergence behave like almost periodic functions, we know there are at least $\asymp T$ zeros of height up to T also in the region of absolute convergence, distributed approximately periodically (cf. [Bes54]).

That's the most I know how to say about their regularity.

APPENDIX A. CODE

A.1. Gamma Computation.

The header `gamma.hpp`:

```

1 // gamma.hpp
2 #ifndef DLD_GAMMA_HPP
3 #define DLD_GAMMA_HPP
4
5 #include <complex>
6
7 namespace dld {
8     std::complex<double> gamma(const std::complex<double>& z);
9 }
10 #endif

```

And the code, largely as shown above.

```

1 // gamma.cpp
2 #include "gamma.hpp"
3 #include <acb.h>
4 #include <arb.h>
5 #include <arf.h>
6
7 namespace dld {
8     std::complex<double> gamma(const std::complex<double>& z) {
9         const slong prec = 53;
10        acb_t zz, res;
11        acb_init(zz); acb_init(res);
12
13        acb_set_d_d(zz, z.real(), z.imag());
14        acb_gamma(res, zz, prec);
15
16        double real = arf_get_d(
17            arb_midref(acb_realref(res)),
18            ARF_RND_NEAR
19        );
20        double imag = arf_get_d(
21            arb_midref(acb_imagref(res)),
22            ARF_RND_NEAR
23        );
24        acb_clear(zz); acb_clear(res);
25        return std::complex<double>(real, imag);
26    }
27 }

```

A.2. Fibonacci Zeta.

I have a minimal header.

```

1 // fibo.hpp
2 #ifndef DLD_FIBO_HPP
3 #define DLD_FIBO_HPP
4
5 #include <complex>
6
7 std::complex<double> zodd_smart(const std::complex<double> & s);
8
9 #endif

```

The code looks annoying only because C++ can be verbose. It's precisely (1), using standard library computations and the gamma function computation above.

```

1 // fibo.cpp
2 #include "fibo.hpp"
3 #include "gamma.hpp"
4 #include <cmath>
5 #include <complex>
6
7 typedef std::complex<double> complex;
8 const long double PI = 3.141592653589793;
9
10 // log( (1 + sqrt(5))/2 )
11 const double logeps = 0.481211825059603;
12 const double sqrt5 = 2.23606797749979;
13
14 complex compute_argument(const complex & s, int m) {
15     return s/2.0 + complex(0.0, PI * m / (2*logeps) );
16 }
17
18 complex summand(const complex & s, int m) {
19     complex ret;
20     if (m % 2 == 0) { ret = 1.0; }
21     else { ret = -1.0; }
22     ret *= std::gamma(compute_argument(s, m));
23     ret *= std::gamma(compute_argument(s, -m));
24     return ret;
25 }
26
27 complex zodd(const complex & s, int limit=1000) {
28     const double logsqrt5 = std::log(sqrt5);
29     complex ret = std::exp(s * logsqrt5);
30     ret /= (8.0 * std::gamma(s) * logeps);
31     complex sum = summand(s, 0);
32     for (int m = 1; m < limit; m++) {
33         sum += 2.0 * summand(s, m);
34     }
35     return ret * sum;
36 }
37
38 complex zodd_smart(const complex & s) {
39     int limit = static_cast<int>(std::ceil(std::abs(s.imag())));
40     if (limit < 20) { limit = 20; }
41     return zodd(s, 2*limit);
42 }

```

A.3. Making the grid. I use two versions of this, depending on whether I'm making an argument plot or a contour plot. The commented out lines indicate the other version.

```

1 #include "fibo.hpp"
2 #include <iostream>
3 #include <fstream>
4 #include <complex>

```

```

5  #include <thread>
6  #include <vector>
7  #include <string>
8
9  void write_grid_chunk(int start_row, int end_row, int grid_size,
10    double min_xval, double max_xval,
11    double min_yval, double max_yval,
12    const std::string& chunk_filename) {
13    std::ofstream file(chunk_filename);
14
15    // file << "real,imag,arg\n";
16    file << "real,imag,rpart,ipart\n";
17    // Iterate over the assigned grid chunk
18    for (int i = start_row; i < end_row; ++i) {
19      for (int j = 0; j < grid_size; ++j) {
20        double real_part = min_xval + (max_xval - min_xval) * i / (grid_size - 1);
21        double imag_part = min_yval + (max_yval - min_yval) * j / (grid_size - 1);
22        std::complex<double> s(real_part, imag_part);
23        std::complex<double> result = zodd_smart(s);
24        // double arg_normalized = (std::arg(result) + PI) / (2 * PI);
25        // file << real_part << "," << imag_part << "," << arg << "\n";
26        double rpart = std::real(result);
27        double ipart = std::imag(result);
28        file << real_part << "," << imag_part << "," << rpart << "," << ipart << "\n";
29    }
30  }
31 }
32
33 // Function to combine all chunk files into the final output file
34 void combine_files(const std::string& output_filename, int num_chunks) {
35   std::ofstream outfile(output_filename);
36   std::string chunk_filename;
37   // Write the header
38   // outfile << "real,imag,arg\n";
39   outfile << "real,imag,rpart,ipart\n";
40
41   // Combine the chunk files
42   for (int i = 0; i < num_chunks; ++i) {
43     chunk_filename = "args_chunk_" + std::to_string(i) + ".csv";
44     std::ifstream infile(chunk_filename);
45     std::string line;
46     // Skip the header line of the chunk file
47     std::getline(infile, line);
48     // Copy the rest of the chunk file to the final output file
49     while (std::getline(infile, line)) {
50       outfile << line << "\n";
51     }
52     infile.close();
53   }
54   outfile.close();
55 }
```

```

57 void save_arggrid_to_file(int grid_size,
58   double min_xval, double max_xval,
59   double min_yval, double max_yval,
60   const std::string& filename) {
61   int num_threads = 6;
62   // Divide the grid rows by the number of threads
63   int chunk_size = grid_size / num_threads;
64   std::vector<std::thread> threads;
65   for (int t = 0; t < num_threads; ++t) {
66     int start_row = t * chunk_size;
67     int end_row = (t == num_threads - 1) ? grid_size : (t + 1) * chunk_size;
68     std::string chunk_filename = "args_chunk_" + std::to_string(t) + ".csv";
69
70     // Start a new thread to write a portion of the grid
71     threads.push_back(
72       std::thread(write_grid_chunk, start_row, end_row,
73       grid_size, min_xval, max_xval, min_yval, max_yval,
74       chunk_filename)
75     );
76   }
77
78   for (auto& t : threads) {
79     t.join();
80   }
81
82   combine_files(filename, num_threads);
83
84   // Clean up chunk files
85   for (int i = 0; i < num_threads; ++i) {
86     std::string chunk_filename = "args_chunk_" + std::to_string(i) + ".csv";
87     std::remove(chunk_filename.c_str());
88   }
89
90   std::cout << "Data saved to " << filename << std::endl;
91 }
92
93 int main() {
94   int grid_size = 400;
95   double min_xval = -35.0;
96   double max_xval = 5.0;
97   double min_yval = -5.0;
98   double max_yval = 35.0;
99   std::string filename = "sizes.csv";
100  save_arggrid_to_file(grid_size, min_xval, max_xval, min_yval, max_yval, filename);
101  return 0;
102 }

```

A.4. Plotting Code.

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3 import csv
4

```

```

5  def plot_contours(filename):
6      x_vals = []
7      y_vals = []
8      r_vals = []
9      i_vals = []
10     with open(filename, 'r') as csvfile:
11         reader = csv.DictReader(csvfile)
12         for row in reader:
13             x_vals.append(float(row['real']))
14             y_vals.append(float(row['imag']))
15             r_vals.append(float(row['rpart']))
16             i_vals.append(float(row['ipart']))
17     # Convert to numpy arrays
18     x_vals = np.array(x_vals)
19     y_vals = np.array(y_vals)
20     r_vals = np.array(r_vals)
21     i_vals = np.array(i_vals)
22
23     # Reshape arg_vals into grid
24     grid_size = int(len(r_vals)**.5 + 0.5)
25     r_grid = r_vals.reshape((grid_size, grid_size), order='F')
26     i_grid = i_vals.reshape((grid_size, grid_size), order='F')
27
28     plt.figure(figsize=[10, 10])
29     plt.contour(i_grid, colors="#f9ae54",
30                 extent=[x_vals.min(), x_vals.max(), y_vals.min(), y_vals.max()],
31                 origin="lower")
32     plt.contour(r_grid, colors="#0482d7",
33                 extent=[x_vals.min(), x_vals.max(), y_vals.min(), y_vals.max()],
34                 origin="lower")
35     plt.tight_layout()
36     plt.savefig("plot.png")
37
38 plot_contours('sizes.csv')

```

REFERENCES

- [AKLDW24a] Eran Assaf, Chan Ieong Kuan, David Lowry-Duda, and Alexander Walker. The Fibonacci zeta function and continuation, 2024. Forthcoming. (Cited on page 1)
- [AKLDW24b] Eran Assaf, Chan Ieong Kuan, David Lowry-Duda, and Alexander Walker. The Fibonacci zeta function and modular forms, 2024. Forthcoming. (Cited on page 1)
- [Bes54] Abram Samoilovitch Besicovitch. *Almost periodic functions*. Dover New York, 1954. (Cited on page 6)
- [Joh17] F. Johansson. Arb: efficient arbitrary-precision midpoint-radius interval arithmetic. *IEEE Transactions on Computers*, 66:1281–1292, 2017. (Cited on page 2)

- [LD20a] David Lowry-Duda. `phase_mag_plot`. https://github.com/davidlowryduda/phase_mag_plot/, September 2020. [Online; Reference version at <https://doi.org/10.5281/zenodo.4035117>]. (Cited on page 2)
- [LD20b] David Lowry-Duda. The Fibonacci zeta function and modular forms. Talk at Dartmouth, available at <https://davidlowryduda.com/notes-from-a-talk-at-dartmouth-on-the-fibonacci-zeta-function/>, May 2020. (Cited on page 1)
- [LD21] David Lowry-Duda. Visualizing modular forms. Talk at Oregon, available at https://davidlowryduda.com/wp-content/uploads/2021/05/visualizing_modular_forms-compressed.pdf, May 2021. (Cited on page 3)
- [Sag20] Sage Developers. *SageMath, the Sage Mathematics Software System (Version 8.8)*, 2020. <https://www.sagemath.org>. (Cited on page 2)