The Polylog and the Riemann Hypothesis

Linas Vepstas

29 October 2023

Abstract

The Riemann Hypothesis can be restated as a hypothesis about the zeros of the polylog function. The relationship is quite simple, clear and direct: the nontrivial zeros of the Riemann zeta are in direct correspondence with the zeros of the polylog. This correspondence allows RH to be refined into two distinct versions: a strong form, where polylog zeros are uniformly bounded away, and a weak form, where RH holds, but flirts with failure.

The breakdown of RH would suggest some pretty weird behaviors in the polylog that would run counter to conventional ideas about holomorphic functions. Perhaps this opens up a new route for searching for RH proofs. Things still look hard: one would have to prove that certain polylog "varieties" are "entire", in a certain sense.

The correspondence between polylog zeros and RH zeros can be clearly seen in animations of the polylog function along the critical strip, posted posted on the linas.org website (See https://linas.org/art-gallery/polylog/polylog.html) This short note is meant to accompany and explain those animations, providing additional detail and explanations not given on the website.

The animations provide insight that is otherwise difficult to obtain: the polylog function has a complicated behavior, due in part to it's having two branch points, and participates in many identities.

Introduction

The polylog function has the series expansion

$$\text{Li}_s(z) = \sum_{n=1}^{\infty} \frac{z^n}{n^s}$$

This converges for $|z| < 1$ and $\Re s > 1$. Analytic continuation can be used to extend this to other values of s and z. Comparing the above to the series representation for the Riemann zeta

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}$$

one is tempted to make a naive (but entirely incorrect) algebraic manipulation that suggests $\zeta(s) = \text{Li}_s(1)$. But this is false. The problem is that the polylogarithm has a branch point (essential singularity) at $z = 1$; the polylog has many sheets, wrapped as a helix around $z = 1$. The naive identity cannot hold.

Although one does not have this simple relationship between the Riemann zeta and the polylogarithm, one still has a behavior that makes it "almost true". Zeros of the polylogarithm "hit" the branch point exactly when s corresponds to a nontrivial Riemann zero. The relation is deeper: it seems that the zeros of the polylogarithm, when viewed on the critical line $s = 1/2 + i\tau$, are in exact correspondence with the nontrivial zeros of the Riemann zeta.

The relationship is best exposed by creating an animation of the polylog along the critical line $s = 1/2 + i\tau$, varying τ over time. Several animations can be found at https://linas.org/art-gallery/polylog/polylog.html. The animations can be explained by starting with a still. The figure below shows the phase of $\mathrm{Li}_s(z)$ on the complex-z plane, for fixed $s = 0.5 + i14$.

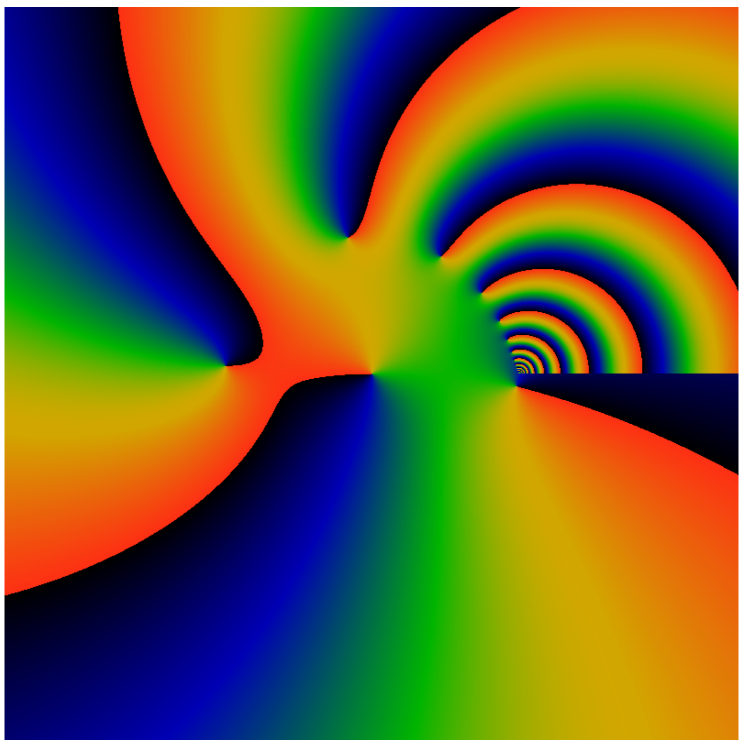

The point $z = 0$ is at the center of the figure; the domain of the figure is over the intervals $-2.5 \le \Re z, \Im z \le 2.5$. The phase[1] is color-coded so that black corresponds to $-\pi$, moving through blue at $-\pi/2$, green at 0, yellow at $+\pi/2$ and red at $+\pi$. Thus, each sharp red–black transition encodes a change of phase by 2π. These edges terminate at zeros of $\mathrm{Li}_s(z)$, where the phases wrap around a point: this is the content

[1]The word "phase" is just the conventional definition. Write $\mathrm{Li}_s(z) = a + ib = M(s;z)\exp i\phi(s;z)$ with both M and ϕ being real functions. Then $M = \sqrt{a^2 + b^2}$ is the magnitude, and $\phi = \arctan b/a$ is the phase. As always, the phase is ambiguous up to a multiple of 2π.

of Cauchy's theorem. Prominently, at the center of the image,[2] we see that $\mathrm{Li}_s(0) = 0$.

The discontinuity extending to the right from the branch point at $z = 1$ is the branch cut. A precise expression for the difference $\Delta = \mathrm{Li}_s(x + i\varepsilon) - \mathrm{Li}_s(x - i\varepsilon)$ across the branch cut is given in my main paper on the polylog;[1] that paper provides both algorithms for high-precision computation, as well as discussions of the monodromy. In particular, there is a second branch point, at $z = 0$, which appears on other sheets. Thus, continuous paths on the complex z plane can wind around either of these two branch points.

The above image shows a sequence of zeros, located near the circle $|z| = 1$. That they are only near the circle, but not exactly on, becomes evident by watching the movies on the web page https://linas.org/art-gallery/polylog/polylog.html. The sequence of zeros on the upper half-plane accumulates onto the branch point at $z = 1$. The precise form of the accumulation is given in the next section. It is exponential in the distance to the branch-point.

Just underneath the branch cut, in the lower half-plane, a lone zero is visible. It is the one terminating the red–black edge, just underneath the branch point. This zero is a zero of the polylogarithm. As s is slowly increased from $s = 1/2 + i14$ to $s = 1/2 + i14.134725$, this polylog zero will smack into the branch point. This is clearly visible in the movies (watch the movies now, if you have not watched them yet.)

What does this mean? The locations of the polylog zeros vary smoothly as a function of s (it cannot be otherwise, the polylog is holomorphic in s.) As one moves along the critical line $s = 1/2 + i\tau$, slowly increasing τ, one discovers that the polylog zeros peel off the branch point, spin round the origin, and return, hitting the branch-point *exactly*, whenever τ is one of the nontrivial zeros of the Riemann zeta. That is, the polylog zeros are Riemann zeros, in the making, yet to be born. They circle about smoothly, and become "true" Riemann zeros when they hit the branch point. Now is a good time to watch the movies, if you haven't seen them. Take particular note of what is happening in the following frames:

14.134725
21.022040
25.010858
30.424876
32.935062
37.586178
40.918719

The above list are the first few Riemann zeros. The movies now provide a basis of discussion for the rest of this text: what does the Riemann hypothesis look like, when re-expressed in terms of the polylog? It is not hard to figure this out, but still, it is entertaining to ponder. Nothing "deep" is happening here; it is, in a sense, "obvious". If this opens up a new way of thinking about a proof for the RH, then the path remains unobvious. I don't know of any easy way of proving theorems about the zeros of the

[2]For $z \ll 1$ small, the first term of the series dominates: $\mathrm{Li}_s(z) \approx z = |z| e^{i\phi}$ and so $\phi \approx 0$ to the right, along the positive real axis (and thus, green), and $\phi \approx -\pi/2$ along the negative imaginary axis (and thus blue), etc. The red/black edge is just the $\phi \approx \pm\pi$ along the negative real axis.

polylog. Nonetheless, a number of interesting statements can be made. Most interesting perhaps is that RH can be refined into a strong and a weak version, only one of which may hold.

Accumulation Point

The picture above clearly shows an accumulation point of polylog zeros, accumulating onto the branch point at $z=1$. They seem to sit, at least approximately, on the unit circle. In the animations, they appear to spread out as τ increases. This section derives an explicit formula for their approximate location; it matches this observed behavior.

The most important observation here is that the zeros occur in a sequence. They can be unambiguously labeled by an integer n, counting up from one. They come in sequential order, so each precedes the last. In the animations, it can be seen that the zeros, once they've circled the origin, play a game of "Flying Dutchman", and interleave themselves back into the sequence. This eventual interleaving does not alter the order of their "birth" from the accumulation point. This remains sequential and unambiguous. This sequencing is one of the important properties of the polylog zeros.

The accumulation of zeros near the branch point can be explored by working with Tom Apostol's "periodic zeta function" $F(q;s)=\mathrm{Li}_s\left(e^{2\pi iq}\right)$. This function is defined in Apostol's textbook.[2] The analytic behavior, as $q\to 0$, is obtained in my polylog paper.[1] The periodic zeta is first rewritten in terms of the Hurwitz zeta

$$\zeta(s,q)=\sum_{n=0}^{\infty}\frac{1}{(n+q)^s}$$

The precise equivalence is

$$\zeta(1-s,q)=\frac{\Gamma(s)}{(2\pi)^s}\left[e^{-i\pi s/2}F(q;s)+e^{i\pi s/2}F(1-q;s)\right]$$

Writing $s=\sigma+i\tau$ and setting $1\ll\tau$ kills the $e^{i\pi s/2}$ term, leaving

$$\begin{aligned}F(q;s)=&\frac{(2\pi)^s}{\Gamma(s)}e^{-i\pi s/2}\zeta(1-s,q)+\mathscr{O}\left(e^{-\pi\tau/2}\right)\\ \approx&K\zeta(1-s,q)\end{aligned}$$

for some constant K independent of q. For $\sigma=1/2$, it can be shown that $|K|=1$, i.e. that it is a pure phase. This happens because $\Gamma(1/2+i\tau)\sim e^{\pi\tau/2}$ for $1\ll\tau$, killing the other factors.

As one approaches $z=1$, one has that $q\to 0$ and the first term in the Hurwitz series summation dominates the others. Thus, one has

$$\zeta(s,q)=q^{-s}+C+\mathscr{O}(q)$$

for some constant C that depends on s but is independent of q. Writing $q=re^{i\phi}$ and assuming $r\ll 1$, this expands to

$$\zeta(s,q)=r^{\phi\tau-\sigma}e^{-i(\phi\sigma+\tau\log r)}+C+\mathscr{O}(r)$$

Combining the above, we deduce that

$$\begin{aligned} F(q;s) =& K\left(C+q^{s-1}\right)+\mathscr{O}(q) \\ =& K\left(r^{1+\sigma-\phi\tau}e^{i(\phi\sigma+\tau\log r)}+C\right)+\mathscr{O}(r) \end{aligned}$$

Thus, one finds $F(q;s)=0$ when $r^{1+\sigma-\phi\tau}=|C|$ and $\phi\sigma+\tau\log r=\text{const.}-2\pi n$ for positive integer n and a constant of $\mathscr{O}(1)$. Both can be simultaneously solved. For $1\ll\tau$, we can take $\phi\approx 0$, bumping it only enough to get the magnitude correct. For the phase, setting $\phi\sigma\approx 0$ yields a sequence of zeros, located at

$$r_n\approx e^{\text{const.}/\tau}\times e^{-2\pi n/\tau}$$

This demonstrates the observed accumulation point. As $n\to\infty$ we get a sequence $r_n\to 0$. Since $\phi\approx 0$, we have that q is approximately real, $q_n\approx r_n$ and so the zeros line up quite near the unit circle as $n\to\infty$. This formula also predicts increasing spacing of the zeros, as τ increases. This is what is seen in the movies.

Note that the above has also generated a sequence numbering for the zeros. The numbering might not start at precisely $n=1$, due to the constant term that was not assiduously tracked. Thus, the first few zeros have to be labeled "by hand", while the remaining zeros occur near q_n in sequential order. The message here is that the zeros can be explicitly labeled, in sequential order, as they are sprouted from the branch point.

Riemann zeta zeros

It seems that the the polylog zeros can be placed in correspondence with the nontrivial Riemann zeta zeros. This section sketches the process.

Lets recap some salient facts. First, as s is varied by small amounts, the zeros of $\text{Li}_s(z)$ move about smoothly. There can be no other way: $\text{Li}_s(z)$ is analytic in s; a small change in s just shifts the zeros smoothly (and analytically). In particular, the zeros are "conserved": they cannot pop into and out of existence as s is changed. With one exception: they can pop out of branch points, or disappear into them. At all other locations, they move about smoothly. (This is basic complex analysis and won't be belabored here.)

Although the intro noted that $\zeta(s)\neq\text{Li}_s(1)$, a glance at the image above indicates that the $z\to 1$ limit is path dependent. The periodic zeta shows how to do it: avoid the singularity at $q\to 0$ by taking $q\to 1$ instead. Actually, any $q\to N$ for any integer $N>0$ should work. Explicitly,

$$\lim_{q\to N}\sum_{n=1}^{\infty}\frac{e^{2\pi iq}}{n^s}=\zeta(s)$$

and this time, the direction of approach should not matter.[3]

[3] I hate being pedantic, but I know I have readers who get upset when I'm not. So, independent of

These two ingredients are sufficient to explain the behavior of the polylog zeros as τ is varied on the critical line, i.e. on $s = 1/2 + i\tau$. That is to say, as τ is increased (as in the movies), each of the polylog zeros rotate around, counterclockwise, and hit $q = 1$ exactly when τ is a nontrivial Riemann zero. There are now several questions that arise:

1. (Circulation.) As τ increases, it appears that the polylog zeros rotate (circulate) around the origin. Does this circulation hold indefinitely, as τ grows?

2. (Correspondence.) For the case of $s = 1/2 + i\tau$, do any of the polylog zeros fail to hit $q = 1$? That is, as they rotate around, do they ever pass to the left or right of the branch-point?

3. (RH.) If one performs the same animations along $s = \sigma + i\tau$ with $\sigma \neq 1/2$, do any of the zeros hit $q = 1$? If they did, this would be a violation of RH.

4. (Bracketing.) If $\sigma < 1/2$, the movies show that the polylog zeros cross the real number line to the left of $z = 1$. Can it ever happen that they would cross to the right? Conversely, if $1/2 < \sigma$, it seems that every polylog zero crosses to the right of $z = 1$. In the movie, it passes through the cut onto the next sheet. Does this bracketing hold for all polylog zeros?

Lets explore these a bit more. For question 1, the movies suggest that circulation should continue indefinitely. It is hard to imagine how it might break down. Not clear how to prove this, but it should be noted that the periodic zeta is oscillatory in $|q|$. Each oscillation appears to be associated with a zero, and the period of oscillation increases as q gets large. In the previous section, the oscillations of the periodic zeta were explicitly linked to polylog zeros, as $q \to 0$. The same link appears to also hold as $|q| \to \infty$. Perhaps with a bit of work, this could be used to prove the circulation hypothesized in question 1.

If the second question can be answered in the negative (and if RH holds), then we have that the polylog zeros can be put in correspondence with the nontrivial Riemann zeros. That is, each and every polylog zero eventually becomes a Riemann zero.

Careful observation of the movies indicates that each polylog zero continues to circle around, slotting itself back into the sequence ("playing a game of Flying Dutchman") and becoming a Riemann zero again, a second, third, forth ... time as well. Even as new zeros are sprouted from the branch point, the old ones get recycled. Overall, it would seem that the number of circulating zeros becomes more and more dense.

It would presumably be quite interesting to know the interleaving sequence. It seems reasonable to think that perhaps this interleaving sequence corresponds to the orbit of some geodesic on some peculiar Riemann surface. This seems to be a recurring theme in the industry, but is a bit out-of-scope for the current paper.

---

approach means that one writes $q = N + \varepsilon \exp i\phi$ with both ε and ϕ being real. One takes the limit $\varepsilon \to 0$ and demonstrates uniform convergence: for all ϕ there exists an upper bound, and convergence is better than this upper bound. For $N > 0$ this is always possible. Well, that, and $\Re s > 1$, else analytically continue to remove pole at $s = 1$, etc. which does complicate the eqns but is entirely tractable. Again, see my polylog paper.[1]

The Riemann Hypothesis

The third question is identical to the Riemann Hypothesis (RH): If the answer is "yes", then there exists a polylog zero s such that $\zeta(s) = 0$ but $\sigma \neq 1/2$. Conversely, if the answer is "no", then the RH is satisfied (all nontrivial zeros lie on the critical line.)

Consider the case of a movie with $\sigma \neq 1/2$. Assuming the circulation hypothesis (given above,) each polylog zero will eventually loop about, and cross the real line. That is, z will be pure-real when it crosses, and, in order to preserve the RH, it must have $z \neq 1$. This suggests that it should be possible to bound each crossing to occur at some distance δ away from $z = 1$. Lets refine the definition of bracketing, and convert the third question into a delta-epsilon hypothesis that's equivalent to the RH.

Write $\sigma = 1/2 + \varepsilon$ for ε real. Let m be an integer counting the m'th polylog zero to rotate through (past) the real axis. Let τ_m be the value of τ when the m'th crossing happens. Likewise, $s_m = \sigma + i\tau_m$. Basically, we're applying labels to each crossing. Let x_m be a real number, the location of the crossing polylog zero on the real axis. That is, $\mathrm{Li}_{s_m}(x_m) = 0$. Write $\delta_m = |1 - x_m|$ as the distance of the crossing point from $z = 1$.

Assuming RH, we then have that, for each $|\varepsilon| > 0$ and each polylog zero m, we must have $\delta_m > 0$. Clearly, the converse holds as well: if there's always such a $\delta_m > 0$, then RH is true. So, this holds if and only if.

There are two possibilities for the behavior of the δ_m if RH is true. One is that they are uniformly bounded away from $z = 1$. That is, for each $|\varepsilon| > 0$ there exists a uniform bound $\delta > 0$ such that $\delta_m \geq \delta$ for all m. The other possibility is that there is no uniform bound, and that $\delta_m \to 0$ as $m \to \infty$. The first possibility could be called "strong RH" or "uniform RH". The second could be called "weak RH". In a certain sense, "strong RH" says that RH never even gets close to failing, while "weak RH" says that RH is true, but just barely, flirting with disaster. The word "weak" is a fun choice; it is reminiscent of weak convergence in Hilbert spaces.

I have no idea if anyone has explored this before, or anything equivalent to this. Write me, let me know.

The Polylog Variety

The polylog variety can be defined as the collection $V \subset U(\mathbb{C}) \times \mathbb{C}$ where the polylog vanishes:

$$V = \{(z, s) \in U(\mathbb{C}) \times \mathbb{C} \,|\, \mathrm{Li}_s(z) = 0\}$$

The space $U(\mathbb{C})$ is the covering space for the complex holomorphic structure of $\mathrm{Li}_s(z)$. It stands for $\mathbb{C}^\infty$ modulo the polylog monodromy: it is the collection of sheets, glued together at the branch points, such that $\mathrm{Li}_s(z)$ is "entire", after excluding the branch points. No such cover is needed for the s coordinate: there is a simple pole at $s = 1$; there are no branch points.

The shape of the variety V can be understood as a collection of two-dimensional sheets. From the work on the accumulation point, above, we know that a polylog zero $\mathrm{Li}_s(z) = 0$ can be labeled with an integer m, at least, when s is held fixed. That is, $z = z_m$. The location of z_m in the complex-z plane, or, more precisely, its location in the cover $U(\mathbb{C})$, is a smooth function of s. That is, $z_m = z_m(s)$. The movies show what happens as both σ and τ are varied, each as the other is held fixed. For a fixed m, the zero $\text{Li}_s(z_m(s)) = 0$ defines a smooth mapping $z_m : \mathbb{C} \to U(\mathbb{C})$. Based on the labeling given by the accumulation point, it appears that there are countably many such sheets.

There is also a corresponding variety generated by the poles of the polylog. These are not visible on the principal sheet, but a clockwise rotation about the $z = 1$ branch-point brings these into view. These are shown in one of the movies. These not far away from the zeros in the sheet below; they move in a similar manner, and thus form a collection $p_m : \mathbb{C} \to U(\mathbb{C})$ where $\text{Li}_s(p_m(s)) = \pm\infty$ denotes the poles.

The movies show that, when holding τ fixed, and moving towards $\sigma \to \pm\infty$, one watches z_m flee either to infinity, or drop to zero. In the other direction, holding σ fixed and varying $\tau \to \pm\infty$, one has the budding-then-circulating motion for each z_m, as described earlier. Very approximately, $z_m(s)$ behaves like $e^{-2\pi m/s}$, up to assorted corrections. This includes the "correction" that $z_m(s)$ circulates to other sheets, when $\sigma > 1/2$.

Ribbons

If RH is violated, then the bracketing hypothesis is also violated (since the corresponding polylog zero fails to bracket.) It is well-known that that the RH is violated, then the violating zero must be doubled. That is, for real $\varepsilon > 0$, if it happens that $\zeta(1/2 + \varepsilon + i\tau) = 0$ then one also has that $\zeta(1/2 - \varepsilon + i\tau) = 0$. This mirroring follows from the reflection formula for ζ. From question three, above, we have that there must be some polylog zero, call it z_m, which corresponds to this RH violation. As τ passes τ_m, then z_m passes through the $q = 1$ point. More precisely, this happens in both the $1/2 \pm \varepsilon + i\tau$ movies. What happens to z_m in the $1/2 \pm \alpha + i\tau$ movies, where $-\varepsilon < \alpha < \varepsilon$? Clearly, each of those z_m must circulate past the the real axis as well. Each of these z_m is a smooth function of α, so that $z_m(\alpha)$ is a smooth curve (and therefore continuous). Thus, when $\alpha = 0$,the corresponding z_m also passes $q = 1$ either at the same time τ_m, or possibly a bit earlier or later. If it misses $q = 1$ entirely, then there has to be another pair of $\pm\alpha$ on the curve that do hit $q = 1$, and the reasoning can be repeated, this time with a tighter bound. At any rate, if there is a violation, then z_m is associated not just with a pair of mirrored RH zeros, but also with a third zero that is on the critical line. There could also be five, seven or any odd number of zeros associated with z_m. (I hope I've explained all this clearly.)

One way to envision the process above is to imagine the variety $z_m(s)$ described in the previous section. If RH is true, then we expect only one point s in all of $z_m(s)$ to correspond to a nontrivial Riemann zero. This is the point $s_m = 1/2 + i\tau_m$ corresponding to the m'th nontrivial Riemann zero. Sheets and zeros are in one-to-one correspondence.

If RH is violated, then a sheet m will have multiple points s for which $\zeta(s) = 0$ (and always an odd number, per the argument above). This suggests a perhaps novel approach for proving RH: demonstrate that all of the functions $z_m(s)$ are one-to-one; that is, demonstrate that $z_m(s) \neq z_m(s')$ whenever $s \neq s'$. This would hold if one could show that $z_m(s)$ are "entire". I'm using scare-quotes, because the range of z_m is not all of $\mathbb{C}$, but the covering space $U(\mathbb{C})$. This seems to add a lot of baggage. Perhaps, though, it is not entirely impossible. Holomorphic functions also have this desired property, at least in limited domains. And it seems $z_m(s)$ is holomorphic. If one can prove that $z_m(s)$ has no zeros or poles, then one would be done. And a casual examination does suggest that $z_m(s)$ has no zeros or poles, and thus entire. This would constitute a proof of RH. Actually arriving at the required details to explicitly demonstrate all of these claims seems daunting. It seems one would need assorted general theorems on varieties, together with specializations for the polylog, to achieve this in any kind of manageable manner.

In short, the bracketing hypothesis (property?) works with the smoothness of the locations of the polylog zeros to force fairly strong constraints on what can happen if RH is violated. If RH is violated, then there are these "violation arcs" associated with the violating z_m. If not, then there aren't any such arcs: they're all shrunken to a single point.

Conclusion

We've demonstrated two different hypothesis, named "strong" and "weak", both of which imply RH. Conversely, RH implies that either "strong" or "weak" must hold.

That's all for now.

References

[1] Linas Vepstas, "An efficient algorithm for accelerating the convergence of oscillatory series, useful for computing the polylogarithm and Hurwitz zeta functions", *Numerical Algorithms*, 47(3), 2008, pp. 211–252, preprint: arXiv:math.CA/0702243 (February 2007).

[2] Tom M. Apostol, *Introduction to Analytic Number Theory*, Springer; New York, 1976.