A CERTIFIED UPPER BOUND FOR THE DE BRUIJN–NEWMAN CONSTANT

KENAN

Abstract. Let H_t denote the de Bruijn–Newman heat deformation of the Riemann ξ-function in the normalization of D. H. J. Polymath, and let Λ be the associated de Bruijn–Newman constant. We give a computer-assisted proof that

$$\Lambda \leq \frac{1729}{10000} = 0.1729.$$

We apply the Polymath15 upper-bound criterion. Its integer parameter is $X = 6000000185827$; the other parameters are $t_0 = 3377/20000$ and $y_0 = 9/100$. Its three inputs are: the published verification of the Riemann hypothesis through height 3000175332800; an interval-arithmetic argument-principle certificate on an enlarged intermediate barrier; and a cutoff-uniform final-time canopy combining 1,756 finite cells with two analytic tails. The finite checker uses a sparse four-prime Euler mollifier. For every output coefficient it evaluates the state at the initial cutoff and after every reachable divisor threshold, thereby avoiding a nonmonotonic terminal-cutoff substitution. A complete regression checks all 451 reachable zero-through-four-prime state patterns and rejects an omitted-initial-state counterfeit invisible to bounded small-cutoff testing. The result improves the established unconditional bound 11/50 by exactly 471/10000. It is a fixed-architecture certified upper bound, not an optimality theorem.

1. Introduction

Let

$$H_t(z) = \int_0^\infty e^{tu^2}\Phi(u)\cos(zu)\ \mathrm{d}u$$

be the classical heat deformation of the Riemann ξ-function. The functions H_t are even entire functions, real on the real axis, and satisfy the backward heat equation. The work of de Bruijn and Newman shows that there is a constant Λ such that H_t has only real zeros precisely for $t \geq \Lambda$ [1, 3]. The Riemann hypothesis is equivalent to $\Lambda \leq 0$, while Rodgers and Tao proved the complementary inequality $\Lambda \geq 0$ [6].

D. H. J. Polymath developed an effective Riemann–Siegel approximation for H_t, together with a criterion that converts three explicit zero-free regions into an upper bound for Λ [5]. Their unconditional implementation gave

$$\Lambda \leq \frac{11}{50} = 0.22.$$

The purpose of this paper is to certify a smaller parameter pair in the same theorem:

$$t_0 = \frac{3377}{20000} = 0.16885, \qquad y_0 = \frac{9}{100} = 0.09.$$

Date: July 22, 2026.

2020 *Mathematics Subject Classification.* Primary 11M26; Secondary 30D15, 65G20.

Key words and phrases. De Bruijn–Newman constant, Riemann xi-function, heat flow, interval arithmetic, certified computation, argument principle.

Theorem 1.1. *In the Polymath15 normalization of the de Bruijn–Newman heat family,*

$$\boxed{\Lambda \le \frac{1729}{10000} = 0.1729.}$$

Consequently the established unconditional upper bound $11/50$ *improves by the exact amount*

$$\frac{11}{50} - \frac{1729}{10000} = \frac{471}{10000} = 0.0471.$$

The numerical size of the improvement is not the only issue. A finite canopy checker used in an earlier certificate evaluated signed Euler-mollified coefficients at the terminal Riemann–Siegel cutoff while claiming every intermediate integer cutoff. The relevant triangle functional is not monotone under activation of a new divisor term. The repair used here evaluates every reachable coefficient state. A second failure mode is subtler: a regression over small cutoffs can agree with the literal state machine on its whole test domain while omitting a pattern that occurs at the theorem scale. We therefore enumerate the complete abstract state space for zero through four Euler primes.

The proof has four layers. Section 2 records the exact imported criterion and parameters. Section 3 closes the initial-time region using the published Platt–Trudgian finite-height theorem. Sections 4 and 5 certify the effective-model transfer and the intermediate barrier. Section 6 proves the final-time ray by a cutoff-uniform finite canopy and complementary analytic tails. Section 9 describes the proof boundary and the executable artifact.

2. The upper-bound criterion

We quote the criterion in the form used by the certificate.

Proposition 2.1 (Polymath15 upper-bound criterion)**.** *The following is Theorem 1.2 of* [5]*. Suppose* $t_0, X > 0$ *and* $0 < y_0 \le 1$*. Assume:*

(i) $\zeta(\sigma + iT) \ne 0$ *for*

$$\frac{1+y_0}{2} \le \sigma \le 1, \qquad 0 \le T \le \frac{X}{2};$$

(ii) $H_{t_0}(x+iy) \ne 0$ *for*

$$x \ge X + \sqrt{1-y_0^2}, \qquad y_0 \le y \le \sqrt{1-2t_0};$$

(iii) $H_t(x+iy) \ne 0$ *on the corresponding intermediate barrier.*

Then

$$\Lambda \le t_0 + \frac{1}{2}y_0^2.$$

The intermediate barrier may be replaced by the larger box

$$X \le x \le X+1, \qquad y_0 \le y \le 1, \qquad 0 \le t \le t_0.$$

We take

$$X = 6000000185827, \qquad t_0 = \frac{3377}{20000}, \qquad y_0 = \frac{9}{100}. \tag{1}$$

Exact rational arithmetic gives

$$t_0 + \frac{1}{2}y_0^2 = \frac{3377}{20000} + \frac{81}{20000} = \frac{1729}{10000}. \tag{2}$$

The parameter certificate also proves

$$\sqrt{1-y_0^2} < 1, \qquad \sqrt{1-2t_0} < 0.82,$$

and verifies that the Riemann–Siegel truncation index is constant at $N = 690988$ throughout the enlarged unit-width barrier.

3. The initial-time region

Platt and Trudgian proved that every nontrivial zero of ζ through

$$T_{\mathrm{PT}} = 3000175332800$$

lies on the critical line [4]. Their exact endpoint exceeds the height required in Proposition 2.1:

$$\frac{X}{2} = 3000000092913.5, \qquad T_{\mathrm{PT}} - \frac{X}{2} = 175239886.5.$$

Moreover,

$$\frac{1+y_0}{2} = \frac{109}{200} > \frac{1}{2}.$$

Thus the published theorem excludes every nontrivial zero in the initial-time strip. This use requires critical-line location only; it does not require simplicity.

The same input closes the exact $t = 0$ face of the enlarged barrier. In the chosen normalization,

$$H_0(z) = \frac{1}{8}\xi\left(\frac{1}{2} + \frac{iz}{2}\right).$$

A zero at $z = x + iy$ would, after applying the functional equation and conjugation, produce a nontrivial zeta zero at height $x/2$ and real part $(1+y)/2$. Since the enlarged face reaches only

$$\frac{X+1}{2} = 3000000092914,$$

the remaining published reserve is 175239886.

Proposition 3.1. *For the parameters* (1)*, hypothesis* (i) *of Proposition 2.1 holds, and the complete $t = 0$ face of the enlarged intermediate barrier is zero-free.*

4. Effective-model transfer

Polymath15 Theorem 1.3 writes the exact heat-flow quotient as

$$\frac{H_t(x+iy)}{B_t(x+iy)} = f_t(x+iy) + \mathcal{E}_t(x,y), \tag{3}$$

where B_t is nonzero on the regions used here and f_t is an effective Riemann–Siegel model. All numerical theorem gates are evaluated using Arb midpoint-radius interval arithmetic [2]. The interval program `error_transfer_01729.c` proves uniformly that

$$|\mathcal{E}_t(x,y)| \le 0.000656886754168108994802390099525 < 0.001.$$

The final-time finite canopy uses a simplified exponent and a sparse Euler mollifier. Including the exponent shift, the finite κ-omission, and the four-prime mollifier norm gives total transfer cost δ_{tr} at most

$$0.00113438368318475277316736553698. \tag{4}$$

Every finite cell is required to exceed the raw distance threshold 0.00114. Hence the exact quotient retains reserve

$$0.00114 - \delta_{\mathrm{tr}} \geq 5.616316815247226832634463022 14 \cdot 10^{-6} > 0. \tag{5}$$

The same verifier proves the monotonicity properties used to propagate the error bounds along the relevant x- and y-ranges.

5. The intermediate barrier

The stored-sum certificate

```
candidate-X6000000185827-tmax0.1809-d20-v3.arbcert
```

covers the stronger time range $0 \leq t \leq 0.1809$. Its reader checks the exact metadata, matrix dimensions, serialized interval values, Taylor-tail enclosure, constant Riemann–Siegel index, finiteness, and end of file. Trailing data and over-coverage requests are rejected.

On the enlarged barrier through $t = t_0$, the argument-principle verifier uses 844 adaptive rectangles. On each rectangle it proves that the interval polygon avoids zero, that every argument increment avoids the logarithm branch cut, and that the exact quotient in (3) has the same winding number as the effective model. The complete winding enclosure is

$$[-2.87 \cdot 10^{-15}, 2.87 \cdot 10^{-15}].$$

Since the winding number is integral, it is zero. The least global exact-model boundary margin is

$$0.95116094446269871778\ldots > 0.$$

Together with Proposition 3.1, the argument principle proves the intermediate hypothesis.

Proposition 5.1. *For the parameters* (1),

$$H_t(x + iy) \neq 0$$

throughout

$$X \leq x \leq X + 1, \qquad y_0 \leq y \leq 1, \qquad 0 \leq t \leq t_0.$$

6. The cutoff-uniform final-time canopy

At $t = t_0$, set

$$N(x) = \left\lfloor \sqrt{\frac{x}{4\pi} + \frac{t_0}{16}} \right\rfloor.$$

The finite proof begins at $N = 690988$, before the final-time ray starts. The issue is to prove a uniform lower bound on every common integer cutoff inside each finite N-cell.

6.1. **The event-state theorem.** Let $\mathcal{D}$ be the squarefree divisor support of a finite Euler mollifier. At output index r, its coefficient at cutoff N has the form

$$C_r(N) = \sum_{\substack{d\in\mathcal{D}\\ d|r\\ r/d\le N}} \lambda_d u_{r,d}. \tag{6}$$

Because the signs λ_d alternate, the triangle functional used by the canopy need not be monotone when a new divisor term activates. A terminal state at N_2 therefore need not majorize all $N \in [N_1, N_2]$.

Lemma 6.1 (Cutoff-event states)**.** *For fixed r, the coefficient states in $N_1 \le N \le N_2$ are exactly the state at N_1, followed by the states reached when N crosses a divisor threshold r/d in $(N_1, N_2]$. If T_r is the improved-triangle contribution, then*

$$\max_{N_1\le N\le N_2} T_r(C_r(N)) = \max_{s\in\mathcal{S}_r} T_r(s),$$

where $\mathcal{S}_r$ is this finite event-state set. Consequently,

$$\sum_r T_r(C_r(N)) \le \sum_r \max_{s\in\mathcal{S}_r} T_r(s)$$

for every common cutoff $N \in [N_1, N_2]$.

Proof. The active divisor set changes only at the finitely many integer thresholds r/d. It is constant between consecutive thresholds, which proves the first identity. The second inequality follows term by term. □

The final inequality is conservative because different output indices may maximize at different common cutoffs, but it is uniform. The need for this repair is witnessed at

$$t = \frac{17}{100}, \quad y = \frac{1}{10}, \quad N = 690988, \quad N_2 = 720000, \quad r = 690990.$$

There the literal intermediate state contributes

$$2.2544469154219889232 \cdot 10^{-7},$$

whereas the terminal state contributes only

$$1.1060396971077443443 \cdot 10^{-7}.$$

Thus terminal substitution underestimates the contribution by more than $1.1484 \cdot 10^{-7}$.

6.2. **Complete pattern-space regression.** Testing Lemma 6.1 only for small N_1, N_2 is not a global state-machine proof. For three Euler primes, the domain $1 \le N_1 \le N_2 \le 25$ realizes only 93 of the 104 reachable three-prime patterns. In particular, on the actual earlier theorem cell

$$N_1 = 690988, \qquad N_2 = 720000, \qquad r = 690990,$$

the state sequence is $(254, 255)$, which does not occur in that small domain. A mutation dropping the initial state 254 passes every bounded case.

The present package directly includes the production sparse checker and constructs every contiguous interval of every divisor-support chain for zero through four primes. It certifies

all zero-through-four-prime patterns	451,
four-prime patterns	353.

On the actual four-prime cell

$$N_1 = 690988, \qquad N_2 = 692000, \qquad r = 691110,$$

the literal sequence is $(32767, 65535)$. A counterfeit dropping the full-support initial state passes the bounded $N_2 \le 24$ regression but is rejected by the complete pattern gate and the actual cell. Terminal-only, endpoint-only, and fabricated-intermediate state sets are also rejected.

6.3. **Finite grids and analytic tails.** The low-y grid has 1,649 cells and covers

$$690988 \le N \le 103999999, \qquad 0.09 \le y \le 0.1809.$$

It uses four Euler primes through $N = 770000$, three through $N = 2000000$, and two thereafter. The high-y grid has 107 cells and covers

$$690988 \le N \le 11999999, \qquad 0.1809 \le y \le 0.82.$$

Every cell is evaluated with interval arithmetic and the event-state algorithm. The weakest raw endpoints are

$$\begin{aligned} \text{low-}y\text{:} \quad & 0.0016881143766071396568, \\ \text{high-}y\text{:} \quad & 0.0241907009721247452 81. \end{aligned}$$

Both exceed 0.00114, and (5) transfers them to strict nonvanishing of the exact quotient.

The complete local arm64 evidence and a decorrelated x86 replay with FLINT 3.6.0 both cover all 1,756 cells. The x86 campaign initially timed out on eight late low-y cells under a 360-second cell bound; all eight were rerun under a 900-second bound and passed.

The complementary analytic tails prove

$$\begin{aligned} &N \ge 104000000, \quad 0.09 \le y \le 1, \quad |f_{t_0}(x+iy)| \ge 0.010729153849097086\ldots, \\ &N \ge 12000000, \quad 0.1809 \le y \le 1, \quad |f_{t_0}(x+iy)| \ge 0.012533669720697907\ldots. \end{aligned}$$

The immediate predecessor starts $N = 103000000$ and $N = 11000000$ fail the strict 0.01 tail gate. Thus the finite-to-tail splice is declared at the first certified starts rather than inferred from nearby values.

Proposition 6.2. *For the parameters* (1),

$$H_{t_0}(x+iy) \neq 0$$

whenever

$$x \ge X + \sqrt{1-y_0^2}, \qquad y_0 \le y \le \sqrt{1-2t_0}.$$

7. Proof of the main theorem

Proof of Theorem 1.1. Proposition 3.1 proves the initial-time zeta hypothesis of Proposition 2.1. Proposition 5.1 proves the enlarged intermediate-barrier hypothesis. Proposition 6.2 proves the final-time ray hypothesis. Applying Proposition 2.1 and (2) gives

$$\Lambda \le t_0 + \frac{1}{2}y_0^2 = \frac{1729}{10000}.$$

□

8. Scope and stopping boundary

The theorem is unconditional, but it is not a proof of the Riemann hypothesis. It is also not an optimality theorem for the numerical architecture. The certificate is tied to the fixed value of X, the Polymath15 effective model, the V3 stored-sum matrix, the sparse finite canopy, the two analytic tails, and the declared interval-arithmetic runtime.

A nonrigorous point scan failed at the nearby target 0.1728. This is a stop for the present raw-threshold route only. It is not evidence that $\Lambda > 0.1728$, and it is not a lower bound of any kind. Improvements may come from a different mollifier, a sharper transfer estimate, a different barrier, a changed X, or another theorem.

The result uses the published finite-height theorem only for zero location. No assertion about zero simplicity enters. Likewise, the complete pattern-space regression proves that the included state enumerator covers the abstract four-prime state space; it does not assert that four primes are universally optimal.

9. Data and code availability

The accompanying reproducibility archive contains the exact theorem statement, proof, dependencies, provenance, source code, certificates, and verifier. It includes:

- the pinned Polymath15 TeX and PDF;
- the sparse four-prime checker and its repair lineage;
- the complete 451-pattern regression and counterfeits;
- the two finite grids and complete local and x86 evidence;
- the parameter, transfer, and analytic-tail programs;
- the V3 matrix and the 844-rectangle barrier transcript; and
- a manifest covering every non-manifest byte.

The default replay authenticates the complete package inventory, checks the published zeta dependency, parses all finite evidence, reruns the parameter, transfer, tail, bounded-enumeration, and complete-pattern programs, tests malformed inputs and state-machine counterfeits, and performs a semantic V3 read. A full option additionally reruns the full barrier and compares its normalized transcript.

The candidate matrix has SHA-256

`cfdacf53752956d288bac174c573b52a6a66fa99b3d74eff8d87a72f0f83b715`.

The full barrier transcript has normalized SHA-256

`eddbc19d764ef45f92a07ffdd5f441634a9622625baa8543848d7f5a5208d622`.

Every numerical statement in the paper is mapped to a machine-readable scope record and to a certificate or source path in the archive. The default replay is designed for a 20-minute wall-clock limit with one worker. The full barrier replay is substantially slower and is separated explicitly from the default verification command.

References

[1] N. G. de Bruijn, *The roots of trigonometric integrals*, Duke Math. J. **17** (1950), no. 3, 197–226, DOI 10.1215/S0012-7094-50-01720-0.

[2] F. Johansson, *Arb: efficient arbitrary-precision midpoint-radius interval arithmetic*, IEEE Trans. Comput. **66** (2017), no. 8, 1281–1292, DOI 10.1109/TC.2017.2690633.

[3] C. M. Newman, *Fourier transforms with only real zeros*, Proc. Amer. Math. Soc. **61** (1976), no. 2, 245–251, DOI 10.1090/S0002-9939-1976-0434982-5.

[4] D. J. Platt and T. S. Trudgian, *The Riemann hypothesis is true up to* $3 \cdot 10^{12}$, Bull. Lond. Math. Soc. **53** (2021), no. 3, 792–797, DOI 10.1112/blms.12460.

[5] D. H. J. Polymath, *Effective approximation of heat flow evolution of the Riemann* ξ *function, and a new upper bound for the de Bruijn–Newman constant*, Res. Math. Sci. **6** (2019), no. 3, article 31, DOI 10.1007/s40687-019-0193-1.

[6] B. Rodgers and T. Tao, *The de Bruijn–Newman constant is non-negative*, Forum Math. Pi **8** (2020), e6, DOI 10.1017/fmp.2020.6.

Independent researcher