The 10^{22}-nd zero of the Riemann zeta function

A. M. Odlyzko

ABSTRACT. Recent and ongoing computations of zeros of the Riemann zeta function are described. They include the computation of 10 billion zeros near zero number 10^{22}. These computations verify the Riemann Hypothesis for those zeros, and provide evidence for additional conjectures that relate these zeros to eigenvalues of random matrices.

1. Introduction

This is a brief report on computations of large numbers of high zeros of the Riemann zeta functions. It provides pointers to sources of more detailed information.

There have been many calculations that verified the Riemann Hypothesis (RH) for initial sets of zeros of the zeta function. The first were undertaken by Riemann himself almost a century and a half ago. Those calculations did not become known to the scientific community until Siegel deciphered Riemann's unpublished notes [**Sie**]. The first published computation, by Gram in 1903, verified that the first 10 zeros of the zeta function are on the critical line. (Gram calculated values for the first 10 zeros accurate to 6 decimal places, and showed that these were the only zeros below height 50. He also produced much less accurate values for the next 5 zeros. See [**Edw**] for more details.) Gram's work was extended by a sequence of other investigators, who were aided by improvements in both hardware and algorithms, with the two contributing about equally to the improvements that have been achieved. The latest published result is that of van de Lune, te Riele, and Winter [**LRW**]. They checked that the first 1.5×10^9 nontrivial zeros all lie on the critical line. Their computations used about 1500 hours on one of the most powerful computers in existence at that time. Since then, better algorithms have been developed, and much more computing power has become available. With some effort at software and at obtaining access to the idle time on a large collection of computers, one could hope to verify the RH for the first 10^{12} zeros in the next year or so, Jan van de Lune has been extending his earlier work with te Riele and Winter, using the algorithms of [**LRW**] and very modest computational resources. By the end of the year 2000, he had checked that the first 5.3×10^9 zeros of the

zeta function lie on the critical line and are simple, even though he was relying on just three not very modern PCs (unpublished).

Starting in the late 1970s, I carried out a series of computations that not only verified that nontrivial zeros lie on the critical line (which was the sole aim of most of the computations, including those of van de Lune, te Riele, and Winter [**LRW**]), but in addition obtained accurate values of those zeros. These calculations were designed to check the Montgomery pair-correlation conjecture [**Mon**], as well as further conjectures that predict that zeros of the zeta function behave like eigenvalues of certain types of random matrices. Instead of starting from the lowest zeros, these computations obtained values of blocks of consecutive zeros high up in the critical strip. The motivation for studying high ranges was to come closer to observing the true asymptotic behavior of the zeta function, which is often approached slowly.

The initial computations, described in [**Od1**], were done on a Cray supercomputer using the standard Riemann-Siegel formula. This formula was invented and implemented by Riemann, but remained unknown to the world until the publication of Siegel's paper [**Sie**]. The highest zeros covered by [**Od1**] were around zero # 10^{12}. Those calculations stimulated the invention, jointly with Arnold Schönhage [**Od2, OS**], of an improved algorithm for computing large sets of zeros. This algorithm, with some technical improvements, was implemented in the late 1980s and used to compute several hundred million zeros at large heights, many near zero # 10^{20}, and some near zero # 2×10^{20}. Implementation details and results are described in [**Od3, Od4**]. These papers have never been published, but have circulated widely.

During the last few years, the algorithms of [**Od3, Od4**] have been ported from Cray supercomputers to Silicon Graphics workstations. They have been used to compute several billion high zeros of the zeta function, and computations are continuing, using spare cycles on machines at AT&T Labs. Some of those zeros are near zero # 10^{22}, and it has been established (not entirely rigorously, though, as is explained in [**Od3, Od4**]) that the imaginary parts of zeros number $10^{22} - 1$, 10^{22}, and $10^{22} + 1$ are

$$1,370,919,909,931,995,308,226.490240...$$
$$1,370,919,909,931,995,308,226.627511...$$
$$1,370,919,909,931,995,308,226.680160...$$

These values and many others can be found at

⟨http://www.research.att.com/∼amo/zeta_tables/index.html⟩.

Further computations are under way and planned for the future. Very soon 10^{10} zeros near zero # 10^{22} will be available. It is likely that some billions of zeros near zero # 10^{23} will also be computed. A revision of [**Od3, Od4**] that describes them is planned for the future [**Od7**]. Results will be available through my home page,

⟨http://www.research.att.com/∼amo⟩.

Finally, let me mention that many other computations of zeros of various zeta and L-functions have been done. Many are referenced in [**Od5**]. There are also interesting new results for other classes of zeta functions in the recent Ph.D. thesis of Michael Rubinstein [**Ru**].

The next section describes briefly the highlights of the recent computations.

2. High zeros and their significance

No counterexamples to the RH have been found so far. Heuristics suggest that if there are counterexamples, then they lie far beyond the range we can reach with currently known algorithms (cf. [Od7]). However, there is still an interest in undertaking additional computations in the ranges we can reach. The main motivation is to obtain further insights into the Hilbert-Pólya conjecture, which predicts that the RH is true because zeros of the zeta function correspond to eigenvalues of a positive operator. When this conjecture was formulated about 80 years ago, it was apparently no more than an inspired guess. Neither Hilbert nor Pólya specified what operator or even what space would be involved in this correspondence. Today, however, that guess is increasingly regarded as wonderfully inspired, and many researchers feel that the most promising approach to proving the RH is through proving some form of the Hilbert-Pólya conjecture. Their confidence is bolstered by several developments subsequent to Hilbert's and Pólya's formulation of their conjecture. There are very suggestive analogies with Selberg zeta functions. There is also the extensive research stimulated by Hugh Montgomery's work on the pair-correlation conjecture for zeros of the zeta function [Mon]. Montgomery's results led to the conjecture that zeta zeros behave asymptotically like eigenvalues of large random matrices from the GUE ensemble that has been studied extensively by mathematical physicists. This was the conjecture that motivated the computations of [Od1, Od3, Od4] as well as those described in this note. Although this conjecture is very speculative, the empirical evidence is overwhelmingly in its favor.

To describe some of the numerical results, we recall standard notation. We consider the nontrivial zeros of the zeta function (i.e., those zeros that lie in the critical strip $0 < \text{Re}(s) < 1$), and let the ones in the upper half of the critical strip be denoted by $\frac{1}{2} + i\gamma_n$, where the γ_n are positive real numbers arranged in increasing order. (We are implicitly assuming the RH here for simplicity. We do not have to consider the zeros in the lower half plane since they are the mirror images of the ones in the upper half plane.) Since spacings between consecutive zeros decrease as one goes up in the critical strip, we consider the normalized spacings

$$(2.1) \qquad \delta_n = (\gamma_{n+1} - \gamma_n) \frac{\log(\gamma_n/(2\pi))}{2\pi} \ .$$

It is known that the average value of the δ_n is 1. The conjecture is that the distribution of the δ_n is asymptotically the same as the Gaudin distribution for GUE matrices.

Figure 1 compares the empirical distribution of δ_n for $1,006,374,896$ zeros of the zeta function starting with zero $\# 13,048,994,265,258,476$ (at height approximately $2.51327412288 \cdot 10^{15}$). The smooth curve is the probability density function for the normalized gaps between consecutive eigenvalues in the GUE ensemble. The scatter plot is the histogram of the δ_n. The point plotted at $(0.525, w)$ means that the probability that δ_n is between 0.5 and 0.55 is w, for example. As we can see, the empirical distribution matches the predicted one closely.

The paper [Od1] presented similar graphs based on the first million zeros, where the agreement was much poorer, as well as on $100,000$ zeros starting at zero $\# 10^{12}$, where the empirical and GUE distributions matched pretty closely. The graphs in [Od3, Od4], based on large sets of zeros as high as zero $\# 10^{20}$ showed far better agreement, even better than that of Figure 1.
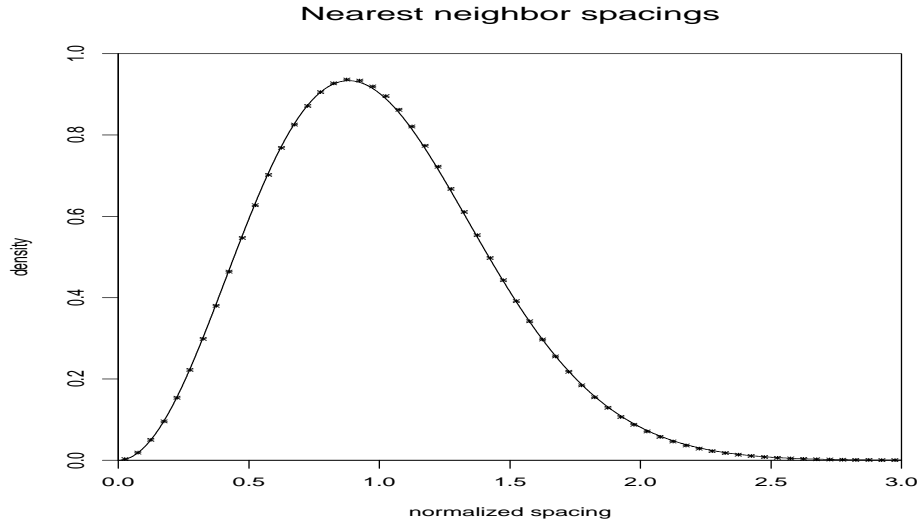
Nearest neighbor spacings

FIGURE 1. Probability density of the normalized spacings δ_n. Solid line: Gue prediction. Scatterplot: empirical data based on a billion zeros near zero # $1.3 \cdot 10^{16}$.

One motivation for continuing the computations is to obtain more detailed pictures of the evolution of the spacing distribution. Graphs such as that of Figure 1 are convincing, but are often inadequate. These graphs do not convey a good quantitative idea of the speed with which the empirical distribution of the δ_n converges to the GUE. They can be misleading, since in the steeply rising parts of the curve, substantial differences can be concealed from the human eye. It is often more valuable to consider graphs such as Figure 2, which shows the difference between the empirical and GUE distributions. This time the bins are of size 0.01, and not the larger 0.05 bins used in Figure 1. It is the large sample size of a billion zeros that allows the use of such small bins, and leads to a picture of a continuous curve. (With small data sets, say of 100,000 data points, which is all that was available in [**Od1**], sampling errors would have obscured what was going on.) Clearly there is structure in this difference graph, and the challenge is to understand where it comes from.

There are many other numerical comparisons between the zeta function and various conjectures that can be performed with large sets of zeros. For example, one can compute moments of the zeta function on the critical line, and compare them with the predictions of the fascinating conjectures of Keating and Snaith [**KeaS**] that relate the behavior of the zeta function at a fixed height to that of eigenvalues of random GUE matrices of a fixed dimension. (The basic Montgomery conjecture only suggested that the asymptotic limits would be the same.) There is also the

FIGURE 2. Probability density of the normalized spacings δ_n. Diffrence between empirical distribution for a billion zeros near zero # $1.3 \cdot 10^{16}$ and the GUE prediction.

general fact that convergence of some properties of the zeta function to asymptotic limits is fast (for example, for the distribution of δ_n), while for others it is slow.

Large scale computations of zeros can also be used in other contexts. In particular, they can be used to improve known bounds on the de Bruijn-Newman constant, as is done in [**Od6**].

Ideally, of course, one would like to use numerical evidence to help in the search for the Hilbert-Pólya operator, and thereby prove the RH. Unfortunately, so far theoretical progress has been limited. Some outstanding results have been obtained, such as the Katz-Sarnak proof that the GUE distribution does apply to zero spacings of zeta functions of function fields [**KatzS1, KatzS2**]. However, these results so far have not been extended to the regular Riemann zeta function.

Acknowledgements. I thank Jeff Lagarias and Jan van de Lune for their comments on earlier versions of this note.

References

[BerK]   M. V. Berry and J. P. Keating, The Riemann zeros and eigenvalue asymptotics, *SIAM Rev.* 41 (1999), 236–266.

[Edw]   H. M. Edwards, *Riemann's Zeta Function*, Academic Press, 1974.

[KatzS1] N. M. Katz and P. Sarnak, *Random Matrices, Frobenius Eigenvalues, and Monodromy*, Colloquium Publications # 45, Amer. Math. Soc., 1999.

[KatzS2] N. M. Katz and P. Sarnak, Zeroes of zeta functions and symmetry, *Bull. Amer. Math. Soc. (N.S.)* 36 (1999), 1–26.

[KeaS]   J. Keating and N. Snaith, Random matrix theory and $\zeta(1/2 + it)$, *Comm. Math. Phys.*
         214 (2000), 57–89.
[LRW]   J. van de Lune, H. J. J. te Riele, and D. T. Winter, On the zeros of the Riemann zeta
         function in the critical strip. IV., *Math. Comp. 46* (1986), 667–681.
[Mon]   H. L. Montgomery, The pair correlation of zeros of the zeta function, pp. 181–193 in
         Analytic Number Theory, H. G. Diamond, ed., Proc. Symp. Pure Math. *24*, Amer. Math.
         Soc., Providence 1973.
[Od1]   A.   M.   Odlyzko,   On   the   distribution   of   spacings   between   zeros   of
         the   zeta   function,   *Math.   Comp.   48*   (1987),   273–308.   Available   at
         ⟨http://www.research.att.com/∼amo/doc/zeta.html⟩.
[Od2]   A. M. Odlyzko, New analytic algorithms in number theory, pp. 466–475 in
         Proc. Intern. Congress Math. 1986, Amer. Math. Soc. 1987. Available at
         ⟨http://www.research.att.com/∼amo/doc/zeta.html⟩.
[Od3]   A.   M.   Odlyzko,   The 10^{20}-th   zero   of   the   Riemann   zeta   function   and
         70   million   of   its   neighbors,   unpublished   manuscript,   1989.   Available   at
         ⟨http://www.research.att.com/∼amo/unpublished/index.html⟩.
[Od4]   A.   M.   Odlyzko,   The 10^{20}-th   zero   of   the   Riemann   zeta   function   and
         175   million   of   its   neighbors,   unpublished   manuscript,   1992.   Available   at
         ⟨http://www.research.att.com/∼amo/unpublished/index.html⟩.
[Od5]   A. M. Odlyzko, Analytic computations in number theory, pp. 451–463 in *Mathe-
         matics of Computation 1943-1993: A Half-Century of Computational Mathe matics*,
         Gautschi, W., ed. Amer. Math. Soc., Proc. Symp. Appl. Math. # 48 (1994). Available
         at ⟨http://www.research.att.com/∼amo/doc/zeta.html⟩.
[Od6]   A. M. Odlyzko, An improved bound for the de Bruijn-Newman constant, *Numerical Al-
         gorithms*, to appear. Available at ⟨http://www.research.att.com/∼amo/doc/zeta.html⟩.
[Od7]   A. M. Odlyzko, *The Riemann Zeta Function: The 10^{22}-nd Zero and Ten Billion of its
         Neighbors*, (tentative title), manuscript in preparation.
[OS]   A. M. Odlyzko and A. Schönhage, Fast algorithms for multiple evaluations of the
         Riemann zeta function, *Trans. Amer. Math. Soc. 309* (1988), 797–809. Available at
         ⟨http://www.research.att.com/∼amo/doc/zeta.html⟩.
[Ru]   M. Rubinstein, Evidence for a spectral interpretation of the zeros of L-functions, Ph.D.
         thesis, Math. Dept., Princeton Univ., 1998.
[Rum]   R. Rumely, Numerical computations concerning the ERH, *Math. Comp.* 61 (1993) 415–
         440, S17–S23.
[Sie]   C. L. Siegel, Über Riemanns Nachlass zur analytischen Zahlentheorie, *Quellen und Stu-
         dien zur Geschichte der Math. Astr. Phys.* 2 (1932), 45–80. Reprinted in *C. L. Siegel,
         Gesammelte Abhandlungen*, Springer, 1966, Vol. 1, pp. 275–310.

AT&T LABS - RESEARCH, FLORHAM PARK, NEW JERSEY 07932
E-mail address: amo@research.att.com