On counterexamples to the Mertens conjecture

Seungki Kim(U. Cincinnati)
Phong Q. Nguyen(ENS/INRIA)

July 19, 2024

Table of contents

- Introduction and summary
- Previous works
- Our method
- Remarks and discussions

The Mertens conjecture

The *Mertens function* is defined as

$$M(x) = \sum_{1 \leq n \leq x} \mu(n),$$

where

$$\mu(n) = \begin{cases} (-1)^k & n \text{ sq.free, has } k \text{ prime factors,} \\ 0 & \text{otherwise} \end{cases}$$

is the Möbius function.

Why does one care about $M(x)$? It is related to the Riemann hypothesis.
More precisely, by elementary manipulations,

$$\frac{1}{\zeta(s)} = s\int_1^\infty \frac{M(x)}{x^{s+1}}dx$$

on some right-half plane, so

$$M(x) = O(x^\theta) \Rightarrow \tfrac{1}{\zeta(s)} \text{ is holomorphic on } \operatorname{Re} s > \theta,$$

i.e., $\zeta(s)$ has no zeros in the region $\operatorname{Re} s > \theta$.

RH is equivalent to saying $\theta < 1/2 + \varepsilon$ for any $\varepsilon > 0$.

The *Mertens conjecture* is a much stronger claim that

$$|M(x)| < \sqrt{x} \text{ for all } x > 1,$$

that dates back to the late 19th century.

Up to 10^{16}, $M(x)$ rarely goes above $0.5\sqrt{x}$.

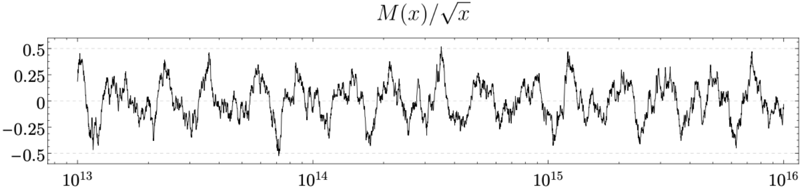

(figure from Hurst (2016))

The Mertens conjecture was disproved in 1985 by Odlyzko-te Riele, in what is perhaps one of the most striking applications of a lattice reduction algorithm to number theory.

Natural follow-up questions:

- What is the smallest value $\mathfrak{x}$ such that $|M(\mathfrak{x})| \geq \sqrt{\mathfrak{x}}$?
- What is a correct asymptotic on $M(x)$?

For the first question, the known conjecture in the literature is

$$\mathfrak{x} \approx \exp(5.15 \times 10^{23})$$

due to Kotnik-van de Lune in 2004, which arises from another conjecture of theirs

$$|M(x)| \approx \frac{1}{2} x^{1/2} \sqrt{\log \log \log x}$$

based on lattice reduction and other numerical methods.

Our result is that the smallest counterexample $\mathfrak{x}$ to the Mertens conjecture is no greater than

$$\approx \exp(1.957 \times 10^{19}),$$

significantly smaller than the Kotnik-van de Lune conjecture. This would impact their conjecture about the growth rate of $M(x)$ as well.

Previous works

To explain our work, we first need to understand the original disproof by Odlyzko-te Riele.

The argument goes by contradiction. Assume the conjecture holds, and for each zero ρ of $\zeta(s)$, associate $\gamma := \mathrm{Im}\,\rho$, $\alpha := |\rho\zeta'(\rho)|^{-1}$, $\psi := \arg(\rho\zeta'(\rho))$. It can be shown that

$$q(x) := \frac{M(x)}{\sqrt{x}} = 2\lim_{n\to\infty} \sum_{\rho:\gamma\in(0,T_n)} \alpha\cos(\gamma y - \psi) + O(x^{-1/2})$$

for some sequence $T_n \sim n$ and $y := \log x$.

(It may help to note $\alpha \to 0$ albeit rather slowly as $\gamma \to \infty$.)

The above suggests a strategy: for some large $T > 0$, find y such that

$$q_T(x) := 2 \sum_{\rho:\gamma\in(0,T)} \alpha \cos(\gamma y - \psi)$$

is large.

Odlyzko-te Riele translates this problem to a case of (simultaneous weighted inhomogeneous) diophantine approximation: provided all $|\gamma y - \psi|_{2\pi}$ are small, we can hope
(here $|x|_{2\pi}$ is the absolute value of x mod $2\pi \in (-\pi, \pi])$

$$q_T(x) = 2 \sum_{\rho:\gamma\in(0,T)} \alpha \cos(\gamma y - \psi) \approx \sum_{\rho:\gamma\in(0,T)} \alpha(2 - |\gamma y - \psi|_{2\pi}^2)$$

to be > 1, since $\sum 2\alpha \to \infty$ as $T \to \infty$.

To find a value of y that makes all $|\gamma y - \psi|_{2\pi}$ small (each weighted by α), Odlyzko-te Riele essentially solves an instance of the *approximate closest vector problem* (aCVP) — the problem of finding a point of a given lattice $L \subseteq \mathbb{R}^n$ reasonably close to a given "target" $\mathbf{t} \in \mathbb{R}^n$.

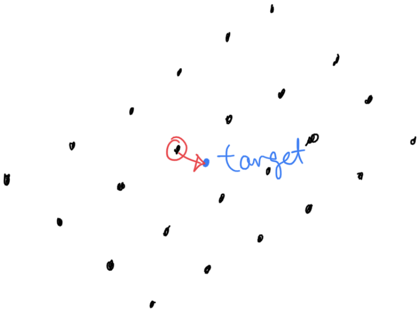

This makes sense, since, for some N,

$(\gamma_1 y - \psi_1, \dots, \gamma_N y - \psi_N)$ small mod 2π

$\Rightarrow (\gamma_1 y - \psi_1 - 2\pi p_1, \dots, \gamma_N y - \psi_N - 2\pi p_N)$ small for some $p_1, \dots, p_N \in \mathbb{Z}$

$\Rightarrow (\gamma_1 y - 2\pi p_1, \dots, \gamma_N y - 2\pi p_1)$ is close to the "target" $(\psi_1, \dots, \psi_N)$

$\Rightarrow (\gamma_1 z 2^{-\mu_1} - 2\pi p_1, \dots, \gamma_N z 2^{-\mu_1} - 2\pi p_1, z 2^{-\mu_2})$ is close to the "target" $(\psi_1, \dots, \psi_N, 0)$, where $z \in \mathbb{Z}$ with $z 2^{-\mu_1} \approx y$, for an appropriate choice of the parameters μ_1, μ_2.

But the former is a point of the lattice generated by rows of

$$\begin{pmatrix} 2\pi & & & \\ & \ddots & & \\ & & 2\pi & \\ \gamma_1 2^{-\mu_1} & \dots & \gamma_N 2^{-\mu_1} & 2^{-\mu_2} \end{pmatrix}.$$

This is essentially the construction of Odlyzko-te Riele, except that extra modifications were made to make the lattice integral, and to weight each $\gamma y - \psi$ according to their coefficients in q_T.

(Some lattice notions I won't explain. . .)

Odlyzko-te Riele uses the *Babai's nearest plane algorithm* to solve the aCVP, which is the standard method to this day.

In order to obtain a high-quality solution, a basis need to be *reduced* i.e., consist of short and orthogonal vectors. Odlyzko-te Riele uses the LLL reduction algorithm, which was state-of-the-art at that time.

A rigorous upper bound on $\mathfrak{x}$ is possible, thanks to

Theorem (Pintz)

Let

$$h_P(y) := 2 \sum_{\gamma < 14000} \alpha \exp(-1.5 \cdot 10^{-6} \gamma^2) \cos(\gamma y - \psi).$$

If there exists $y \in [e^7, e^{50000}]$ *with* $|h_P(y)| > 1 + e^{-40}$, *then* $\mathfrak{x} < \exp(y + \sqrt{y})$.

Again, finding such y comes down to the diophantine approximation problem discussed above.

The value of y found by Odlyzko-te Riele gives $\mathfrak{x} < \exp(3.21 \times 10^{64})$.
Kotnik-te Riele (2006, ANTS XII) improved it to $\mathfrak{x} < \exp(1.59 \times 10^{40})$, by multiple trials over varying parameters.

Recently, Saouter-te Riele (2014) improved Pintz's result as follows.

Theorem (Saouter-te Riele, paraphrased)

Let

$$h_{StR}(y) := 2 \sum_{\gamma < 74000} \alpha \exp(-3 \cdot 10^{-9}\gamma^2) \cos(\gamma y - \psi).$$

If there exists $y \geq 200$ *with* $|h_{StR}(y)| > 1 + 6 \cdot 10^{-8}$, *then* $\mathfrak{x} < \exp(y + \sqrt{y})$.

Using this, and some more trials, they attain $\mathfrak{x} < \exp(1.004 \times 10^{33})$.

All the above results, even the recent ones, use LLL for lattice reduction, which appeared in 1982.

But lattice reduction and other tools has improved drastically since the times of LLL, especially in the last two decades, motivated by post-quantum cryptography. Knowing this, one would naturally try to replace LLL with one of those.

Last year, K. and Rozmarynowycz (then UC undergrad) replaced LLL with BKZ, together with a few additional tricks, which led to $\mathfrak{x} < \exp(1.017 \times 10^{29})$.

The method of the present work

Previous experiences led me to realize a few limitations of the previous approach.

Note that lattice reduction is a rather time-consuming process, especially in three-digit dimensions. So it is infeasible to control all terms of

$$h_{StR}(y) = 2 \sum_{\gamma<74000} \alpha \exp(-3 \cdot 10^{-9}\gamma^2) \cos(\gamma y - \psi).$$

In realistic terms, one is limited to working with the first ≤ 150 terms or so with the largest coefficients.

At first, one may think that reducing the lattice under question really well may yield a winning value of h_{StR} — this takes days to weeks. But this does not yield a particularly impressive value of h_{StR}.

In fact, the quality of aCVP solution and size of h_{StR} correlate somewhat, but not too well, because the "tail" fluctuates by a small yet nontrivial amount.

Instead, one could go for quantity over quality: collect as many candidate points as possible, and hope that in one of those cases the fluctuation occurs in our favor. Luckily, there happens to exist a lattice algorithm that is perfectly suited for this task: *lattice-point enumeration*.

Lattice-point enumeration aims to quickly count all or most of the lattice points within a given ball in the Euclidean space. Roughly, the idea is that it is far more efficient to count those contained within a carefully chosen rectangular/cylindrical shape than those within a literal ball, by using the Gram-Schmidt orthogonalization of a decently well-reduced basis.

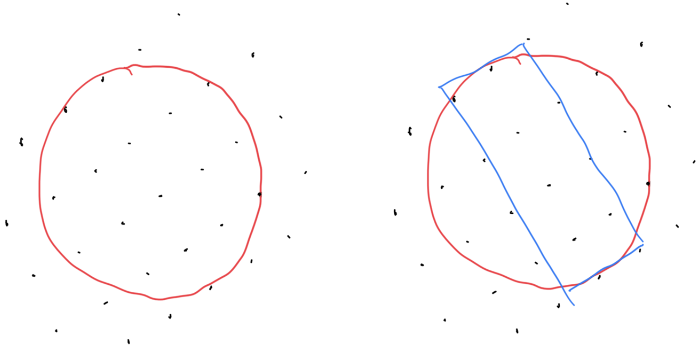

In our work, we used a slight modification of the algorithm of Liu-Nguyen (2013).

It turns out that the lattice under question has a very special shape, almost like $\varepsilon\mathbb{Z} \oplus \mathbb{Z}^N$ for $\varepsilon \approx 10^{-12}$ — nearly orthogonal, with one exceptionally short vector. This leads to a certain amount of time saving:

- On the one hand, this fact itself makes enumeration more efficient than on generic lattices.
- On the other hand, this means along the direction of $\varepsilon\mathbb{Z}$ there are many candidate points with similar values of h_{StR}. So we enumerated on the projection of the lattice onto the orthogonal component of this "short axis," which cut down the search space by a factor of a few million.

Some values we found during a single enumeration in a dimension 141 lattice, among $\sim 17,000$ points found within a few hours of preprocessing + a few hours of counting, on Phong's personal desktop:

y	$h_{StR}(y)$	$y + \sqrt{y}$
889543786 4289868028.044074	-0.974798	8.895×10^{18}
1385953971 0197847064.062257	-0.9949	1.386×10^{19}
1957187885 0562201959.215107	-1.007	1.957×10^{19}
6417170555 7420452732.080835	-1.02	6.417×10^{19}
1 5555848868 6568113612.224656	1.025	1.555×10^{20}
1 8947128314 9477540226.654238	0.997	1.894×10^{20}

From the third line, we obtain our result $\mathfrak{x} < \exp(1.957 \times 10^{19})$.

Some remarks and discussions

- With more time and patience, it may be possible to lower the bound even further.
- With even more time and patience, it may be possible to use our method to numerically investigate the growth rate of $M(x)$. By adjusting parameters, we can heuristically target a specific range of y, enumerate the relevant lattice vectors, and evaluate q_T in that range.

  If we are lucky, this may have a say about the existing conjectures about the growth of $M(x)$:

  $$|M(x)| = O(x^{1/2}(\log\log\log x)^{\theta}),$$

  with $\theta = 1/2$ (Kotnik-van de Lune), $\theta = 1$ (Kaczorowski), $\theta = 5/4$ (Gonek).

- Most importantly, we hope this work helps inform the community of the advances in lattice algorithms, and motivate revisiting other problems where lattice reduction etc. are used in a crucial way e.g., linear relations among the imaginary parts of the zeroes of $\zeta(s)$ — cf. Best-Trudgian (2015).