Lattices and L-functions from nothing

Andrew V. Sutherland

Massachusetts Institute of Technology

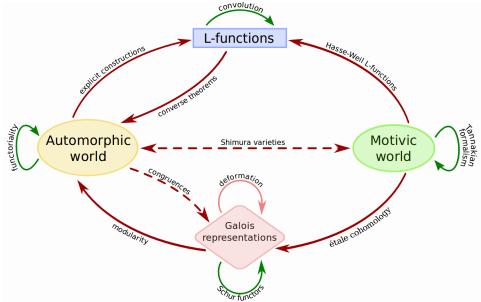

The Simons Collaboration in Arithmetic Geometry, Number Theory, and Computation

(joint work with Andrew R. Booker)

Lattices are ubiquitous in number theory, both as mathematical objects of interest, and as computational tools. I will focus almost exclusively on the latter, which is largely about leveraging efficient methods for lattice reduction (LLL). Examples include:

- xgcd's, matrix kernels, Hermite and Smith normal forms, algebraic relations
- polynomial and ideal reduction, class group and unit group computations
- Diophantine approximation (generalization of continued fractions)
- Baker's method for solving Diophantine equations (linear forms in logarithms)
- Schnorr's algorithm for factoring and discrete logs (finding smooth relations)
- finding rational points on or near an algebraic variety (Elkies' algorithm)
- testing conjectures (ABC, Mertens' conjecture, Hall's conjecture, FLT)
- Coppersmith's method for finding small solutions to polynomial equations

For surveys of these topics, I recommend *The LLL Algorithm: Survey and Applications*, particularly the chapters by Hanrot, May, and Simon.

Factoring integer polynomials in polynomial time

This was the first and most notable application of the LLL algorithm. The timeline:

1967 Berlekamp *Factoring polynomials over finite fields*
Factoring in $\mathbb{F}_p[x]$ is easy for small p (and for large p, *via randomization*).

1969 Zassenhaus *On Hensel factorization, I*
Factoring in $\mathbb{Z}_p[x]$ is easy, which sometimes makes factoring in $\mathbb{Z}[x]$ easy.

1982 Lenstra-Lenstra-Lovász *Factoring polynomials with integer coefficients*
Theorists celebrate (polynomial-time!), practitioners stick with Zassenhaus.

2002 van Hoeij *Factoring polynomials and the knapsack problem*
Practitioners celebrate, theorists scratch their heads (is it even polynomial-time?)

2009 Belebas-van Hoeij-Klüners-Steel *Factoring polynomials over global fields*
Everybody celebrates! (polynomial-time over any global field).

The key to van Hoeij's success was using LLL exactly where it is needed.
We would like to do the same thing with a new algorithm for computing L-functions.

A selective history of L-functions: the early years

1737 Euler proves his Euler product formula: $\sum_n n^{-s} = \prod_p (1 - p^{-s})^{-1}$.

1837 Dirichlet defines Dirichlet L-functions $L(s, \chi) := \sum \chi(n) n^{-s} = \prod (1 - \chi(p) p^{-s})^{-1}$ and uses them to prove his theorem on primes in arithmetic progressions.

1859 Riemann proves that $\zeta(s) := \sum n^{-s}$ has an analytic continuation and satisfies a functional equation, then formulates his Riemann hypothesis (open) and the explicit formula relating zeros of $\zeta(s)$ to prime numbers (all in 8 pages!).

1877 Dedekind introduces Dedekind zeta function $\zeta_K(s) := \sum N(I)^{-s}$.

1916 Ramanujan makes his Ramanujan conjecture (now proved) that the coefficients $\tau(n)$ of the modular form $\Delta(z) = \sum \tau(n) e^{2\pi i n z}$ satisfy $|\tau(p)| < 2p^{11/2}$.

1917 Hecke proves a functional equation for $\zeta_K(s)$ and defines Hecke L-functions.

1922 Mordell proves that for elliptic curves $E/\mathbb{Q}$ the group $E(\mathbb{Q})$ is finitely generated.

1924 Artin introduces Artin L-functions associated to Galois representations, and arising as factors of $\zeta_K(s)$, and conjectures that they are holomorphic (still open).

1933 Hasse proved the Hasse bound $|a_p| \leq 2\sqrt{p}$, where $a_p := p + 1 - \#E(\mathbb{F}_p)$.

A selective history of L-functions: the middle years

1937 Hecke introduces Hecke operators on spaces of modular forms and proves Euler products and functional equations for their cuspidal eigenforms.

1939 Siegel introduces Siegel modular forms and their L-functions.

1939 Rankin and Selberg independently develop Rankin-Selberg L-functions.

1949 Weil formulates the Weil conjectures for smooth projective varieties $X/\mathbb{F}_p$ (now proved) and proves them for curves, including the Weil bound $|a_p| \leq 2g(X)\sqrt{p}$.

1950 Hasse defines the zeta function $\zeta_X(s) = \prod_p \zeta_{X_p}(s) = \zeta(s)\zeta(s-1)/L(X,s)$ of a smooth projective curve $X/\mathbb{Q}$ and formulates the Hasse-Weil conjecture, that the $L(X,s)$ has an analytic continuation and functional equation (open for $g(X) > 1$).

1954 Eichler introduces the Eichler–Shimura construction, associating an isogeny class of elliptic curves $E/\mathbb{Q}$ to each weight-2 newform with integer coefficients.

1955 The modularity conjecture that the Eichler-Shimura construction accounts for all isogeny classes of elliptic curves $E/\mathbb{Q}$ is formulated by Shimura, Taniyama, Weil.

A selective history of L-functions: the middle years

1963 Sato and Tate independently formulate the Sato–Tate conjecture regarding the semi-circular distribution of $a_p(E)/\sqrt{p}$ for typical elliptic curves $E/\mathbb{Q}$.

1960s Birch and Swinnerton-Dyer jointly formulate the BSD conjecture for elliptic curves $E/\mathbb{Q}$, that the rank of $E(\mathbb{Q})$ is the order of vanishing of $L(E,s)$ at $s=1$ (assuming Hasse-Weil), and give an explicit formula for the value of $L^{(r)}(1)$.

1967 In a letter to Weil, Langlands sketches the outlines of the Langlands program, a far-reaching framework of conjectures relating L-functions of automorphic objects (such as modular forms) to those of motivic objects (such as elliptic curves).

1971 Deligne reduces the Ramanujan conjecture to the Weil conjectures (using the Eichler–Shimura correspondence).

1972 Davenport, Swinnerton-Dyer, Tingley, and others tabulate elliptic curves $E/\mathbb{Q}$ of conductor $N \leq 200$, computed their ranks and a_p values for small p, and matched this data to corresponding weight-2 newforms of level N.

1970s Many mathematicians help build the foundations of the Langlands program.

1974 Deligne proves the Weil conjectures in full generality.

A selective history of L-functions: the modern era

1983 Faltings proves Shafarevich's conjecture, implying, in particular, that there are only finitely many abelian varieties $A/\mathbb{Q}$ of a given dimension and conductor.

1980s Frey, Ribet, Mazur, and Serre reduce Fermat's conjecture about integer solutions to $x^n + y^n = z^n$ to the modularity of $E/\mathbb{Q}$ with squarefree conductor.

1989 Kolyvagin proves BSD for modular elliptic curves with analytic rank $r \leq 1$.

1990 Cremona tabulates modular $E/\mathbb{Q}$ up to conductor 1000 (reaches 500000 in 2019).

1995 Taylor and Wiles prove modularity for $E/\mathbb{Q}$ of squarefree conductor.

2000 Breuil, Conrad, Diamond, and Taylor prove modularity for all $E/\mathbb{Q}$.

2011 Taylor and others prove potential modularity results yielding analytic continuation for symmetric power L-functions of $E/\mathbb{Q}$, implying the Sato–Tate conjecture.

2014 Brumer and Kramer formulate the paramodular conjecture, which relates L-functions $L(A, s)$ of typical abelian surfaces $A/\mathbb{Q}$ of conductor N to those of paramodular Siegel newforms of degree 2, weight 2, and level N.

2021 Boxer, Calegari, Gee, Pilloni prove potential modularity for abelian surfaces $A/\mathbb{Q}$.

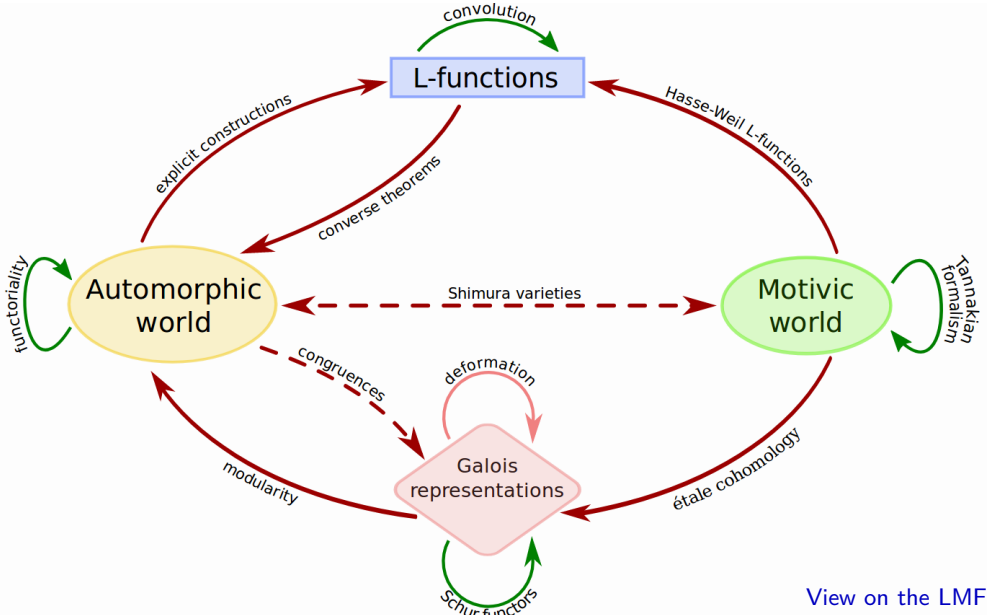

View on the LMFDB

Challenges in moving from dimension one to dimension two

We currently have nothing close to the abelian surface equivalent of the 1972 Antwerp tables of elliptic curves. We know only the first 36 modular abelian surface L-functions unconditionally, of which only 5 are typical (the 1972 Antwerp tables had 749).

- Enumerating paramodular forms of a given level is very difficult; even counting them is hard, due to the absence of dimension formulas. We have provably complete lists of paramodular forms only up to level 353 (all five of them).
- Computing the L-function of a given paramodular form is every difficult; it is usually only feasible to compute a handful of Euler factors.
- There is no analog of the Eichler-Shimura construction for paramodular forms.
- Not all abelian surfaces over $\mathbb{Q}$ are Jacobians of genus 2 curves over $\mathbb{Q}$. One can generically represent an abelian surface as a projective variety in $\mathbb{P}^{15}$ defined by 72 quadratic forms, but this is not a very pleasant thing to do.
- There is no algorithm known to enumerate all genus 2 curves over $\mathbb{Q}$ of a given conductor. Even computing the conductor of a single curve can be very hard.

An axiomatic approach to L-functions (of abelian varieties over $\mathbb{Q}$)

Fix a positive integer g. We shall consider rational L-functions of the form

$$L(s) := \sum_n a_n n^{-s} = \prod_p L_p(p^{-s})^{-1}$$

where

- the a_n are integers that satisfy $|a_n| \leq d_{2g}(n)\sqrt{n}$, where $d_r(n) = \sum_{n_1 \cdots n_r = n} 1$.
- $L(s)$ has an analytic continuation that is holomorphic on $\Re(s) > 0$ for $s \neq 2$.
- $\Lambda(s) := \Gamma_{\mathbb{C}}(s)^g L(s)$ (where $\Gamma_{\mathbb{C}}(s) := 2(2\pi)^{-s}\Gamma(s)$) satisfies a functional equation

$$\Lambda(s) = \varepsilon N^{1-s} \Lambda(2-s)$$

  with conductor $N \in \mathbb{Z}_{>0}$ and root number $\varepsilon = \pm 1$.
- Each $L_p \in \mathbb{Z}[T]$ has $\deg(L_p) \leq 2g$, with equality iff $p \nmid N$.

Under the Hasse–Weil conjecture, every $A/\mathbb{Q}$ of dimension g has such an L-function.

A finite problem

Let $\mathcal{S}(g, N, \varepsilon)$ denote the set of L-functions $L(s)$ that satisfy our axioms for a particular choice of $g, N \in \mathbb{Z}_{>0}$ and $\varepsilon \in \pm 1$.

We expect every $L \in \mathcal{S}(g, N, \varepsilon)$ to be the L-function of a g-dimensional $A/\mathbb{Q}$ (this is far beyond anything we can currently hope to prove, but we don't need to).

Shafarevich's conjecture (proved by Faltings), then implies that $\mathcal{S}(g, N, \varepsilon)$ is finite. Moreover there is an effectively computable $n_0 = O(\sqrt{N})$ for which the coefficients $a_1, \ldots, a_{n_0}$ uniquely determine each $L \in \mathcal{S}(g, N, \varepsilon)$ (and $n_0 = O(\log^2 N)$ under GRH).

We seek an algorithm that takes inputs g, N, ε, determines a suitable n_0, and then outputs a list of distinct tuples $(a_1, \ldots, a_{n_0})$, one for each $L \in \mathcal{S}(g, N, \varepsilon)$.
See Booker and Farmer–Koutsoliotas–Lemurell for prior work in this direction.

Our plan: Compute $\mathcal{S}(g, N, \varepsilon)$ via linear algebra, then search for corresponding $A/\mathbb{Q}$.

Our plan depends crucially on being able to compute $\mathcal{S}(g, N, \varepsilon)$ exactly.
This not only tells us when to stop searching, knowing $a_1, \ldots, a_{n_0}$ helps us search.

The approximate functional equation

Fix g, N, ε. For each nonnegative integer k we define $S_k(x) := \sum_n f_k(n/x) a_n/n$, where

$$f_k(x) := \frac{1}{2\pi i} \int_{c-i\infty}^{c+i\infty} (s-1)^k \Gamma_{\mathbb{C}}(s)^g x^{1-s}\, ds.$$

The functional equation then implies the identity

$$S_k(x) = \varepsilon(-1)^k S_k(N/x),$$

valid for all $x > 0$; this is the *approximate functional equation*. If we choose k so that $(-1)^k = -\varepsilon$ and put $x = \sqrt{N}$ we obtain a nontrivial linear constraint on the a_n:

$$\sum_n f_k(n/\sqrt{N}) a_n/n = 0. \tag{1}$$

The $O(\sqrt{n})$ bounds on a_n and rapid decay of $f_k(x)$ allows us to compute an interval $I_{k,m}$ containing the truncated sum in (1) for $n \le m$ that does not depend on the a_n.

A system of linear constraints

For each $k \geq 0$ of the correct parity (meaning $(-1)^k = -\varepsilon$), we have linear constraints

$$\sum_{n \leq m} f_k(n/\sqrt{N}) a_n/n \in I_{k,m}.$$

These become less useful as k grows, so we restrict to $k = O(N^{1/4})$.
We also have the constraints $|a_n| \leq d_{2g}(n)\sqrt{n}$ for $n \geq 1$.

If we further assume that the $L \in \mathcal{S}(g, N, \varepsilon)$ are automorphic (which we do), we can obtain additional constraints by twisting $L(s)$ by a Dirichlet character $\chi\colon \mathbb{Z} \to \mathbb{C}$, equivalently, taking the Rankin-Selberg convolution of $L(s)$ with $L(\chi, s)$.

This generally increases the conductor and widens the corresponding interval $I_{\chi,k,m}$, but for χ of small conductor (relative to m) and small k we obtain useful constraints

$$\sum_{n \leq m} f_k(n/\sqrt{N}) \chi(n) a_n/n \in I_{\chi,k,m}.$$

Solving the system rigorously using the simplex method

The Euler product for $L(s)$ implies that the a_n are determined by the a_q for prime powers $q = p^e$ with $e \leq 2g$. In order to take advantage of this, and to obtain rigorous results using off-the-shelf simplex solvers with fixed precision, we proceed as follows.

Let $q \leq n_0 < m$ be a prime power. Assume we have recursively fixed values for $a_1, \ldots, a_{q-1}$ that we cannot rule out this sequence as a prefix of a feasible solution.

We now apply the simplex method to a system of linear constraints on variables $a_{q'}$, with q' ranging over prime powers $q \leq q' \leq m$, using the objective functions $\pm a_q$.

The dual solution yields a linear combination of constraints we can compute using interval arithmetic. Plugging in bounds on $a_{q'}$ yields an interval I_q containing a_q.

If $I_q \cap \mathbb{Z}$ is empty, then $a_1, \ldots, a_{q-1}$ is not the prefix of any $L \in \mathcal{S}(g, N, \varepsilon)$. Otherwise, for each $a \in I_q$ we add the tuple $(a_1, \ldots, a_{q-1}, a)$ to our list of feasible tuples.

We continue in this fashion until we run out of feasible prefixes or reach $q = n_0$.

A small example

A short proof that the set $\mathcal{S}(1,13,1)$ is empty, which implies that there are no elliptic curves $E/\mathbb{Q}$ of conductor 13 (this only requires Hasse-Weil, not modularity).

```
> LFN(1,13,-1);
[ 1 ], considering a_2 in { -2, -1, 0, 1, 2 }
[ 1 ], possible a_2: {}

There are no degree 2 motivic weight 1 rational L-functions for N=13 and eps=-1

> LFN(1,13,1);
[ 1 ], considering a_2 in { -2, -1, 0, 1, 2 }
[ 1 ], possible a_2: { -2, -1 }
    [ 1, -2 ], considering a_3 in { -3, -2, -1, 0, 1, 2, 3 }
    [ 1, -2 ], possible a_3: {}
    [ 1, -1 ], considering a_3 in { -3, -2, -1, 0, 1, 2, 3 }
    [ 1, -1 ], possible a_3: { -2 }
        [ 1, -1, -2, -1 ], considering a_5 in { -4, -3, -2, -1, 0, 1, 2, 3, 4 }
        [ 1, -1, -2, -1 ], possible a_5: {}

There are no degree 2 motivic weight 1 rational L-functions for N=13 and eps=1.
```

Timings

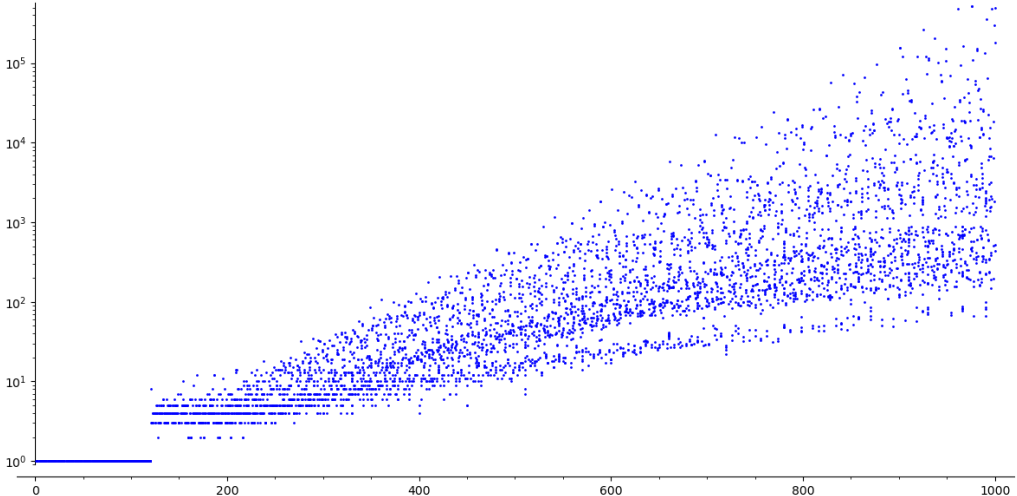

Proving completeness

If our algorithm outputs a nonempty list of feasible tuples $(a_1, \ldots, a_{n_0})$, the next step is to show there is at most one L-function in $\mathcal{S}(g, N, \varepsilon)$ for each prefix.

For this step, if we suppose that $(a_1, \ldots, a_{n_0})$ is the prefix of two distinct L-functions $L(s, \pi_1)$ and $L(s, \pi_2)$ of isobaric cuspidal automorphic representations of $\mathrm{GL}_{2g}(\mathbb{A}_{\mathbb{Q}})$ whose L-functions lie in $\mathcal{S}(g, N, \varepsilon)$. Using the Rankin–Selberg convolution L-function $L(s, \pi_1 \boxtimes \pi_2)$ we construct an inequality which will be violated if n_0 is sufficiently large.

If it is not violated, we increase n_0, extend our tuples, and try again.

Eventually we obtain a list of distinct tuples $(a_1, \ldots, a_{n_0})$, each of which is provably the prefix of at most one automorphic L-function in $\mathcal{S}(g, N, \varepsilon)$.

This gives us an upper bound for our search that we expect to be tight.
Finding an abelian variety for each prefix proves completeness subject to modularity.

We can then attempt to use Faltings-Serre or other methods to prove modularity for each abelian variety. Either our list is complete or we find an explicit nonmodular $A/\mathbb{Q}$.

Searching for genus 2 curves

Over the past several years we have conducted several searches for genus 2 curves of small conductor (we expect to run one more this year). Below is CPU histogram from a computation we ran in 2022 that enumerated more than 10^{19} genus 2 curves using a large parallel computation running on Google cloud platform.

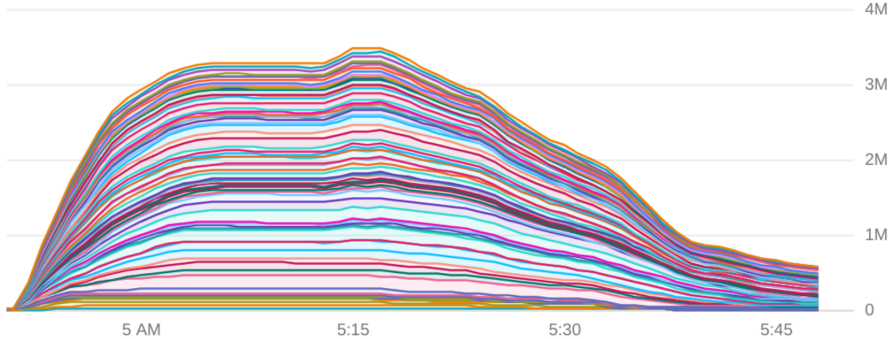

We used a total of 4,034,560 Intel/AMD vCPUs in 73 data centers across the globe.

Searching for genus 2 curves

We found millions of genus 2 curves of small conductor, including the curve

$$C_{903} : y^2 + (x^2 + 1)y = x^5 + 3x^4 - 13x^3 - 25x^2 + 61x - 28$$

of conductor 903 and whose L-function coefficients match those of the paramodular form of level 903 computed by Poor–Yuen that had not previously been matched.

We also found curves of conductor 657, 760, 775, 924 not previously known to occur, and many new genus 2 L-functions of small conductor:

conductor bound	1000	10000	100000	1000000
curves in LMFDB	159	3069	20265	66158
curves found	807	25438	447507	5151208
L-functions in LMFDB	109	2807	19775	65534
L-functions found	200	9409	212890	2426708

A provisional result

Theorem (not yet proved)

Assume the paramodular conjecture.
There are 456 L-functions of abelian surfaces $A/\mathbb{Q}$ with conductor $N \leq 1000$, of which

- *360 arise from products of elliptic curves over $\mathbb{Q}$;*
- *17 arise from weight-2 newforms with quadratic coefficients;*
- *2 arise from the Weil restriction of an elliptic curve over a quadratic field;*
- *77 arise from generic abelian surfaces, of which at least 67 are Jacobians.*

It may be feasible to remove the paramodular hypothesis, but that will depend largely on work by others and it almost certainly won't be feasible much past $N \leq 1000$.

In addition to proving this theorem, we hope to extend it well past $N \leq 1000$.
But this requires algorithmic improvements.

Opportunities for improvement?

Our current approach uses the simplex method in a rather simple-minded way that does not let us exploit integrality to the extent we might like, and uses an objective function that may be suboptimal.

The algorithm often has trouble handling clusters of primes that are close together (e.g. twin primes). Using an objective function that involves more than just the first unknown variable may help, but one would like to optimize the choice systematically.

One possible approach would be to find a small linear combination of constraints whose sum yields a constraint whose coefficients are all close to integers (possibly obtained via LLL). Rounding would yield an integer linear combination of unknown a_q that we could use as an objective function, in the hope of obtaining an interval that contains no integers. This would allow that branch of the computation to terminate.

But there are surely many other ideas worth trying. Please suggest some!

Temporary page!

LaTeX was unable to guess the total number of pages correctly. As there was some unprocessed data that should have been added to the final page this extra page has been added to receive it.
If you rerun the document (without altering it) this surplus page will go away, becau LaTeX now knows how many pages to expect for this document.