Certified Newton schemes for the evaluation of low-genus theta functions

Jean Kieffer (Harvard)

LFANT Seminar, March 22, 2022

Outline

1. Introduction

2. Dupont's algorithm

3. Certified Newton schemes

Theta functions

$g \geq 1$

$\tau \in \mathcal{H}_g$ (i.e. τ is a complex $g \times g$ symmetric matrix and $\operatorname{Im} \tau > 0$)

$z \in \mathbb{C}^g$: column vector

$a, b \in \{0,1\}^g$: theta characteristics.

Theta functions:

$$\theta_{a,b}(z,\tau) = \sum_{m \in \mathbb{Z}^g} E\left((m+\tfrac{a}{2})^t \tau (m+\tfrac{a}{2}) + 2(m+\tfrac{a}{2})^t (z+\tfrac{b}{2})\right)$$

where $E(x) := \exp(i\pi x)$.

Theta constants: value at $z = 0$, as a function of τ.

Why theta functions?

$\theta_{a,b}(z,\tau)$ satisfies many symmetry properties w.r.t. both variables:

- z: quasi-periodic with respect to lattice $\Lambda(\tau)$.
- τ: modular form, i.e. transformation formula under $\mathrm{Sp}_{2g}(\mathbb{Z})$.

They are universal:

Why theta functions?

$\theta_{a,b}(z,\tau)$ satisfies many symmetry properties w.r.t. both variables:

- z: quasi-periodic with respect to lattice $\Lambda(\tau)$.
- τ: modular form, i.e. transformation formula under $\mathrm{Sp}_{2g}(\mathbb{Z})$.

They are universal:

- Lefschetz's theorem: Theta functions (z variable, τ fixed) provide projective embeddings of complex abelian varieties.
- Igusa: Theta functions (z = torsion point, τ variable) realize modular varieties $\Gamma(n^2, 2n^2)\backslash\mathcal{H}_g$ as quasi-projective varieties. Any Siegel modular form can be expressed as a rational fraction in terms of theta functions.

Evaluation of theta functions

Goal

Given (approximations of) $z \in \mathbb{C}^g$ and $\tau \in \mathcal{H}_g$, and given $N \geq 1$, compute approximations of all $\theta_{a,b}(z, \tau) \in \mathbb{C}$ up to an absolute error $\leq 2^{-N}$.

Applications

- CM theory and class polynomials (Enge '09, '14)
- Modular polynomials and isogenies (Enge '09, K. '21)
- Detect subsets of $A[\ell]$ defined over $\mathbb{Q}$...

Works in combination with height bounds/study of denominators.

The naive algorithm

Sum individual terms of the theta series.

$$\theta_{a,b}(z,\tau) = \sum_{m\in\mathbb{Z}^g} E\Big((m+\tfrac{a}{2})^t\tau(m+\tfrac{a}{2}) + 2(m+\tfrac{a}{2})^t(z+\tfrac{b}{2})\Big)$$

- We need all terms $\|m\| \ll \sqrt{N}$, to $\simeq N$ bits of precision.
- E can be computed in quasi-linear time $O(\mathcal{M}(N)\log N)$.
- Total: $O(N^{g/2}\,\mathcal{M}(N)\log N)$.

Lots of possible optimizations.

Uniform in z,τ if suitably reduced.

Main result

Dupont (2006) and Labrande–Thomé (2016) describe a quasi-linear time algorithm in $O_\tau(\mathcal{M}(N)\log N)$ operations to evaluate theta functions. Relies on heuristics.

Main result

Dupont (2006) and Labrande–Thomé (2016) describe a quasi-linear time algorithm in $O_\tau(\mathcal{M}(N)\log N)$ operations to evaluate theta functions. Relies on heuristics.

Theorem (K., 2022)

Variants of Dupont's algorithm yield explicit, provably correct, uniform, quasi-linear time algorithms of cost $O(\mathcal{M}(N)\log N)$ for

- theta functions for $g = 1$
- theta constants for $g = 2$.

In higher genera, we cannot guarantee that the algorithm will work for all (z, τ). If it does, then the output can be certified.

Dupont's algorithm

Theta constants and the AGM (1)

For $\tau \in \mathcal{H}_1$, write

$$\Theta(\tau) = \Big(\theta_{0,0}^2(0,\tau),\ \theta_{0,1}^2(0,\tau)\Big).$$

The duplication formula tells us that

$$\Theta(\tau) \rightsquigarrow \Theta(2\tau)$$

is an AGM step $(x, y) \mapsto (\frac{x+y}{2}, \sqrt{xy})$. There is a sign ambiguity when choosing $\sqrt{x}$ and $\sqrt{y}$.

Theta constants and the AGM (1)

For $\tau \in \mathcal{H}_1$, write

$$\Theta(\tau) = \Big(\theta_{0,0}^2(0,\tau),\ \theta_{0,1}^2(0,\tau)\Big).$$

The duplication formula tells us that

$$\Theta(\tau) \rightsquigarrow \Theta(2\tau)$$

is an AGM step $(x, y) \mapsto (\frac{x+y}{2}, \sqrt{xy})$. There is a sign ambiguity when choosing $\sqrt{x}$ and $\sqrt{y}$.

- Choice of signs is good when $\sqrt{x}, \sqrt{y}$ lie in a quarter plane.
- An AGM sequence with good sign choices converges quadratically to a nonzero value.

Good sign choices

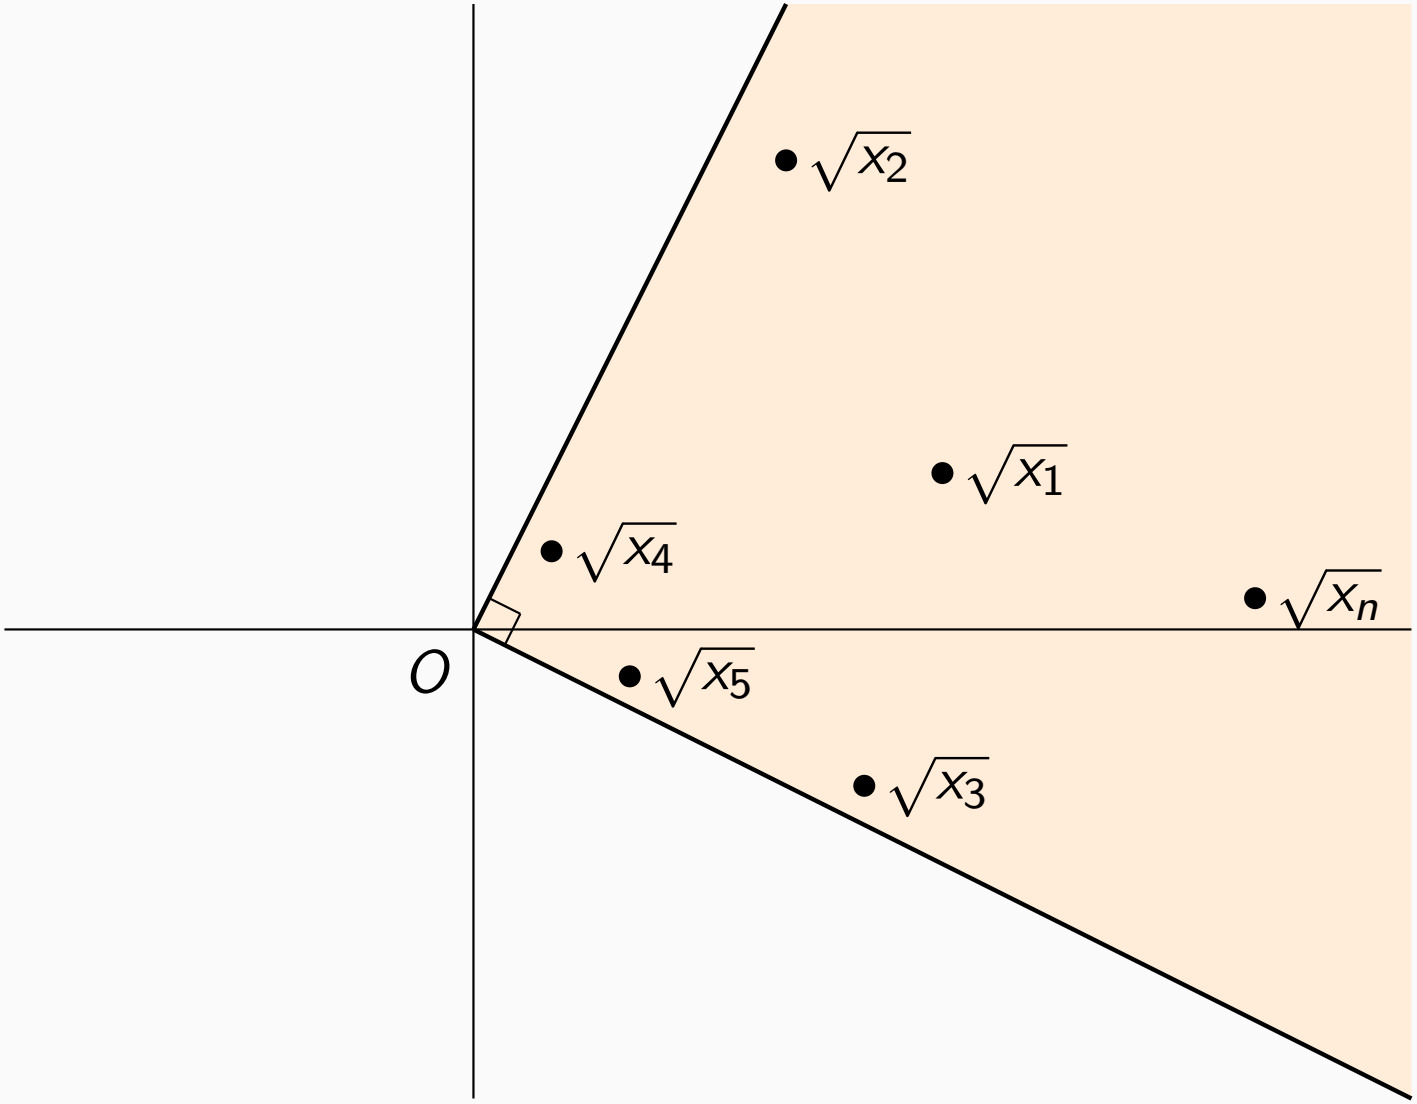

Theta constants and the AGM (2)

We know:

- For each $\tau \in \mathcal{H}_1$, $\left(\Theta(2^n\tau)\right)_{n\geq 0}$ is an AGM sequence.
- Write $q = \exp(i\pi\tau)$, going to zero as $\tau \to \infty$; then

$$\theta_{0,0}(\tau) = 1 + 2q + 2q^4 + O(q^9)$$
$$\theta_{0,1}(\tau) = 1 - 2q + 2q^4 + O(q^9)$$

Consequence: if $\lambda \in \mathbb{C}^\times$, then

$$\Big(\lambda\,\Theta(2^n\tau)\Big)_{n\geq 0}$$

is an AGM sequence and converges quadratically to $\big(\lambda, \lambda\big)$.
We recover $\Theta(\tau)$ without multiplicative factor.

Inversion of theta constants

Use the AGM to invert theta constants.
Input: $\Theta(\tau/2) \in \mathbb{P}^1(\mathbb{C})$.

- Duplication: compute $\left(\theta^2_{a,b}(\tau)\right)_{a,b}$ as a projective point.
- Action by $S = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$: compute $\Theta(S\tau) \in \mathbb{P}^1(\mathbb{C})$ using

$$\theta^2_{0,0}(S\tau) = -i\tau\theta^2_{0,0}(\tau), \qquad \theta^2_{0,1}(S\tau) = -i\tau\theta^2_{1,0}(\tau)$$

  Multiplicative factors cancel.
- Limits of AGM sequences: compute $\theta^2_{0,0}(\tau)$ and $\theta^2_{0,0}(S\tau)$.
  If τ lies in the fundamental domain, all sign choices are good.
- Recover τ using the transformation formula once more.

Complexity: $O_\tau(\mathcal{M}(N)\log N)$.

Dupont's algorithm

Use a Newton scheme.

Given τ, compute $\Theta(\tau/2)$ as follows:

- Compute an approximation Θ_0 of $\Theta(\tau/2)$ at low precision N_0.
- Apply the AGM to compute the corresponding τ_0, close to τ.
- Approximate the derivative D of this AGM function at Θ_0 (finite differences).
- Set $\Theta_1 = \Theta_0 + D^{-1}(\tau - \tau_0)$; it is an approximation of $\Theta(\tau/2)$ to precision $2N_0 - \delta$.
- Repeat until we reach precision N.

Complexity: still quasi-linear time $O_\tau(\mathcal{M}(N)\log N)$.

A uniform algorithm

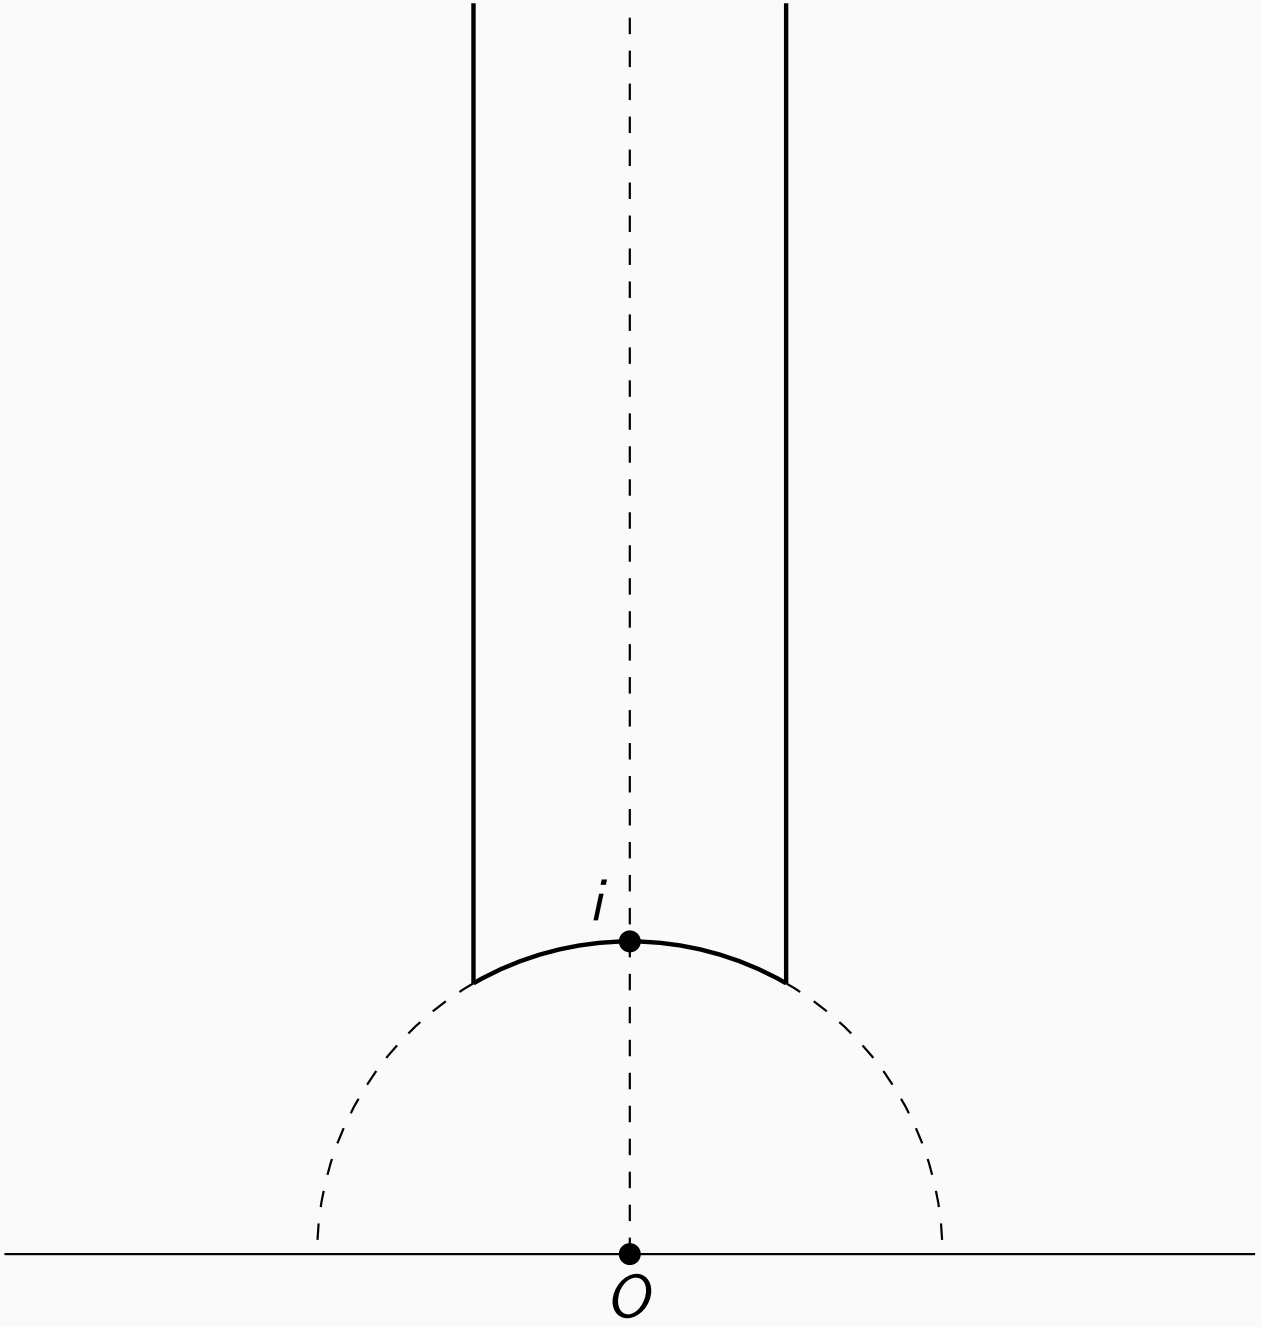

A uniform algorithm

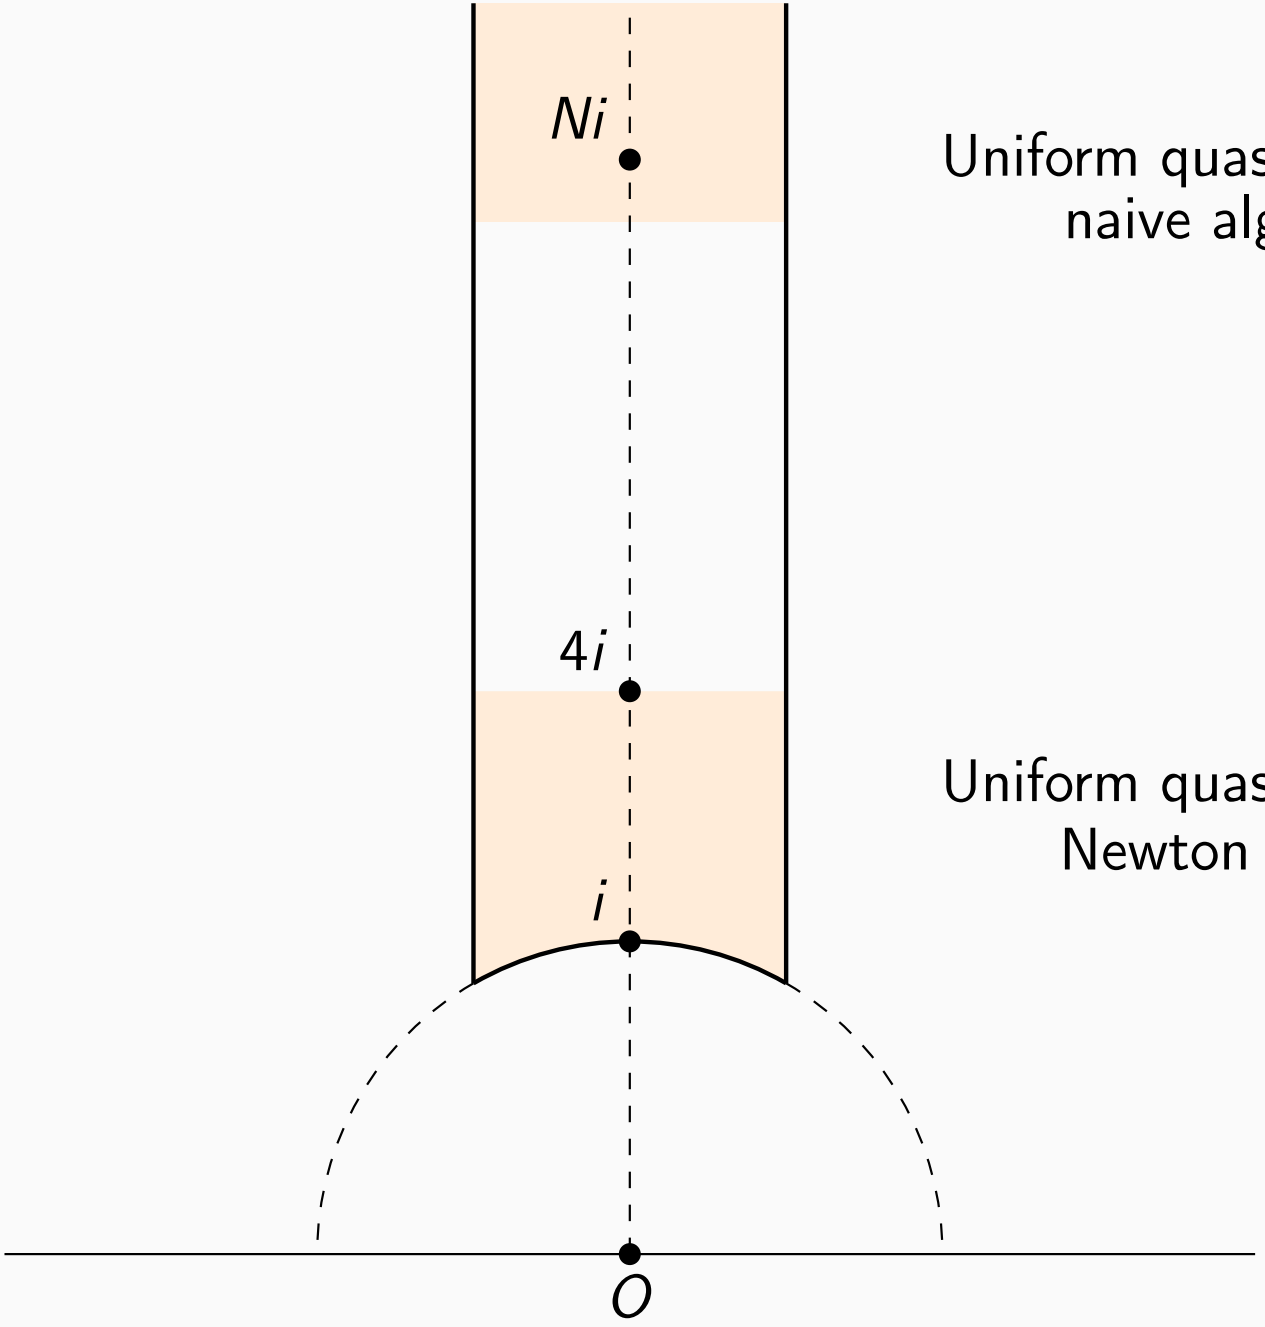

Uniform quasi-linear time
naive algorithm

Uniform quasi-linear time
Newton scheme

A uniform algorithm

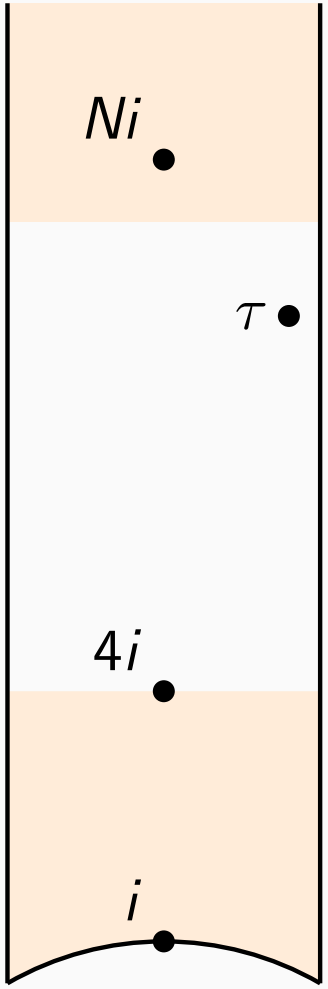

Uniform quasi-linear time
naive algorithm

Uniform quasi-linear time
Newton scheme

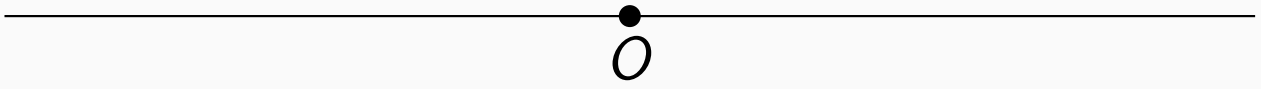

A uniform algorithm

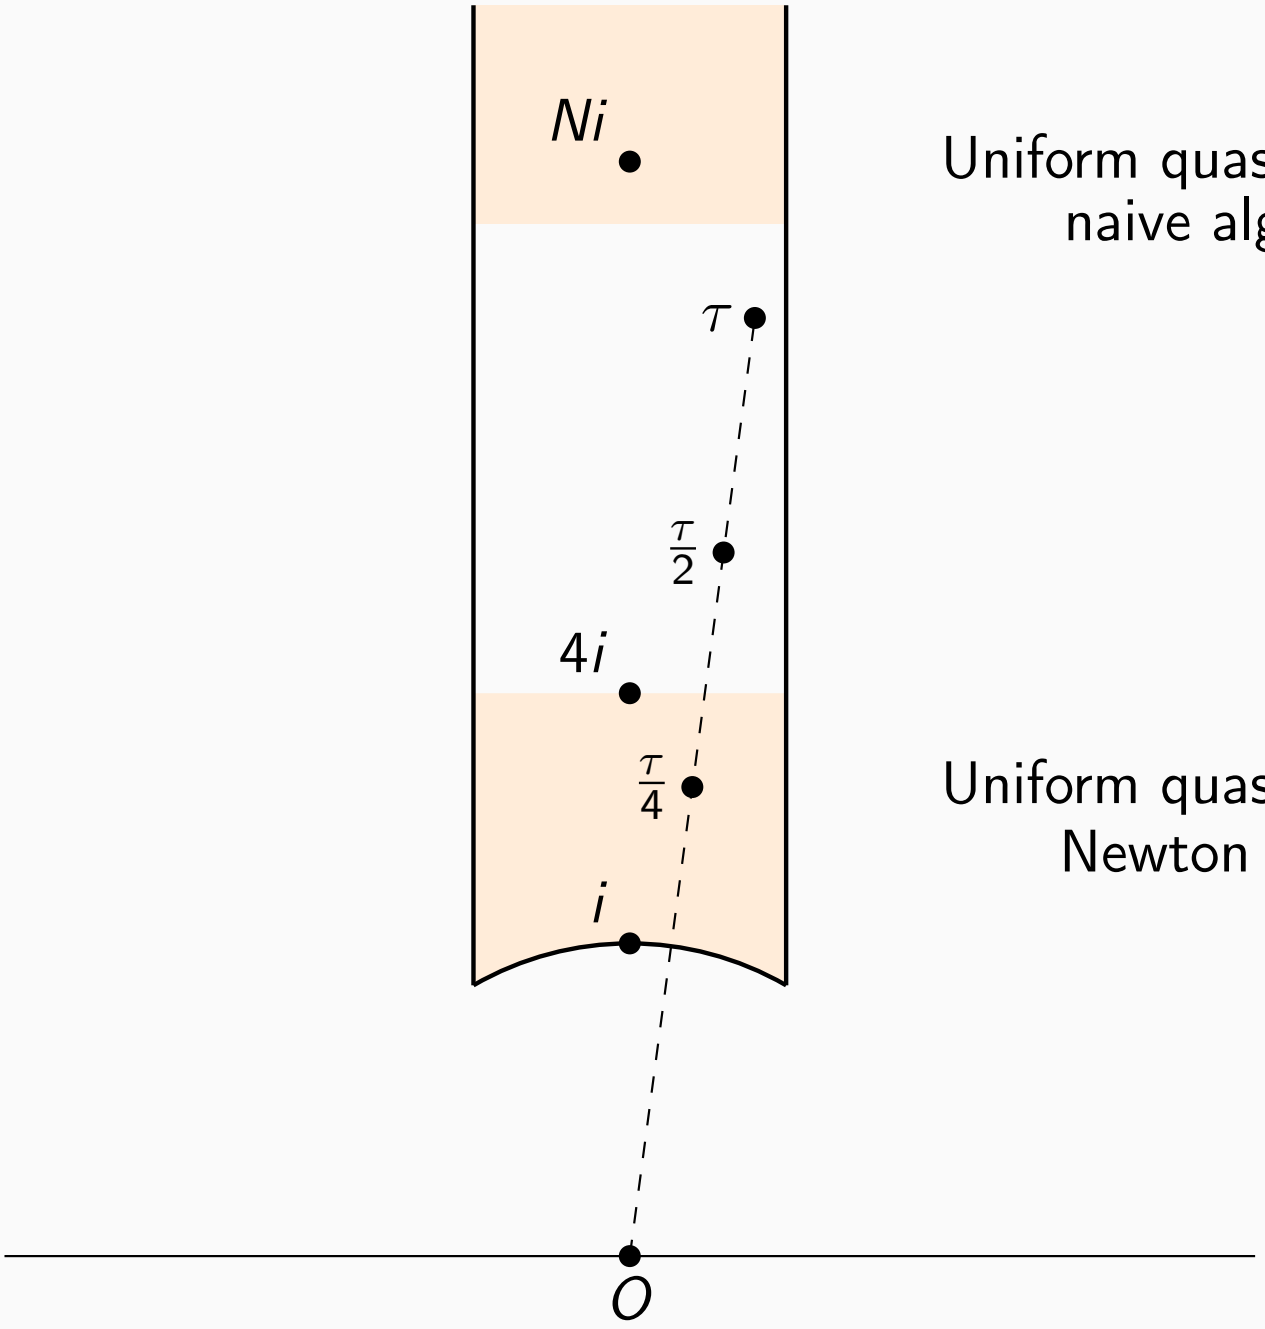

Uniform quasi-linear time
naive algorithm

Uniform quasi-linear time
Newton scheme

Higher genus instances (1)

Dupont's algorithm generalizes to $(z, \tau) \in \mathbb{C}^g \times \mathcal{H}_g$.

- Genus g theta constants ($z = 0$): use

$$\Theta(\tau) = \left(\theta^2_{0,b}(0, \tau)\right)_{b \in \{0,1\}^g}.$$

- Genus g theta functions: use

$$\Theta(\tau) \quad \text{and} \quad \widetilde{\Theta}(z, \tau) := \left(\theta^2_{0,b}(z, \tau)\right)_{b \in \{0,1\}^g}.$$

Dimension of θ-space: $2^g - 1$ or $2^{g+1} - 2$.
Dimension of τ-space: $g(g+1)/2$ or $g(g+3)/2$.

Higher genus instances (2)

- Generalizations of the AGM in higher dimensions: Borchardt sequences. $\Theta(2^n\tau) \to (1, \ldots, 1)$. Similar characterization of quadratic convergence by good sign choices.
- Also consider extended Borchardt sequences (studied by Labrande–Thomé): compute μ from

$$\left(\lambda\Theta(\tau),\, \mu\widetilde{\Theta}(z,\tau)\right)$$

  (Usual Borchardt sequence computes λ.)
- Act by at least $g(g+1)/2$ symplectic matrices $S \in \mathrm{Sp}_{2g}(\mathbb{Z})$. The linearized system should be invertible, in particular square.

For $g = 2$, explicit set of symplectic matrices.

Heuristic aspects

- Describe correct choices of square roots in AGM steps?
- Is the linearized system actually invertible?
- Upper bound on precision loss δ in the Newton scheme?

Make the algorithm uniform in τ?

Certified Newton schemes

Certified Newton schemes

Let $U \subset \mathbb{C}^r$ open, $f : U \to \mathbb{C}^r$ complex-analytic, and $x_0 \in U$.
Let $0 < \rho \leq 1$, $M \geq 1$ and $B \geq 1$ be such that:

- $D(x_0, \rho) \subset U$.
- $\|f(x)\| \leq M$ for each $x \in D(x_0, \rho)$.
- $\|df(x_0)^{-1}\| \leq B$.

Let C be a "nice" function such that $f(x)$ can be evaluated to precision N in time $O(C(N))$ uniformly on $D(x_0, \rho)$.
Then, given

- $f(x_0)$ to precision N,
- x_0 to precision $2\lceil \log_2(2(r+1)M/\rho) \rceil + 2\lceil \log_2(B) \rceil + 4$,

there is an explicit Newton scheme to compute x_0 to precision $N - \lceil \log_2(B) \rceil - 1$ in time $O(C(N))$.

Sketch of proof

Usual explicit bounds for Newton schemes using either:

- Upper bound on $\|df(x_0)\|$, and uniform upper bound on $\|d^2 f(x)\|$ locally around x_0. (Works for $\mathcal{C}^2$ functions.)
- Upper bounds on all derivatives of f at x_0. (Works for real-analytic functions.)

For complex-analytic functions, Cauchy's formula gives both.

Precision losses during the computation can also be managed: Newton schemes have auto-correction.

Limits of Borchardt sequences are analytic (1)

Let $0 < \rho < M$ and

$$U_g(\rho, M) = \left\{ z = (z_i)_{1 \le i \le 2^g} : \rho < \mathrm{Re}(z_i) < M \text{ for all } i \right\}.$$

Theorem

There is a unique analytic function $\lambda \colon U_g(\rho, M) \to \mathbb{C}$ such that $\lambda(z)$ is the limit of the Borchardt sequence with good sign choices starting from z; we have $\rho < \|\lambda(z)\| < M$ on $U_g(\rho, M)$.

Proof

Finite sequences of AGM steps are analytic + locally uniform convergence.

Limits of Borchardt sequences are analytic (2)

Analogous statements for extended Borchardt sequences, but constants are worse.

Without the assumption of good choices of square roots, need to

- Bound the number of bad steps (finite);
- Bound each term away from zero during bad steps.

Good sign choices in low genus

Genus 1 case: Dupont '06 (theta constants), Labrande '18 (theta functions) proved that sign choices are good in all the AGM sequences appearing in Dupont's algorithm, provided that the input is suitably reduced.

Theorem (K. '21)
The same property holds in the case of genus 2 theta constants.

The proof provides explicit lower bounds for the radius of convergence of the AGM functions we are interested in.

This is unlikely to hold verbatim as g grows. However we can still hope for uniform upper bounds on the number of bad steps, etc.

Recall: certified Newton schemes

Let $U \subset \mathbb{C}^r$ open, $f : U \to \mathbb{C}^r$ complex-analytic, and $x_0 \in U$.
Let $0 < \rho \leq 1$, $M \geq 1$ and $B \geq 1$ be such that:

- $D(x_0, \rho) \subset U$.
- $\|f(x)\| \leq M$ for each $x \in D(x_0, \rho)$.
- $\|df(x_0)^{-1}\| \leq B$.

[...]

Invertibility of the linearized system

- If the dimensions of τ-space and θ-space are equal ($g = 1$, genus 2 theta constants):
  The inverse system is entirely described by theta functions.
  We can obtain uniform upper bounds on $\|df^{-1}\|$.
- Higher dimensions: we can either use more symplectic matrices, or equations for the image of θ (which has non-smooth points).
  Obtaining uniform bounds is harder, but we can still certify the Newton scheme independently on each input.

Final computations

Theorem (K. '22)

For suitably reduced input restricted to a compact set, Dupont's algorithm converges in a certified way starting from approximations of at precision

- 60 for genus 1 theta constants,
- 300 for genus 2 theta constants,
- 1600 for genus 1 theta functions.

These are below the practical thresholds with the naive algorithm.

Thank you!

Questions?