Constructing abelian varieties for cryptographic use

Peter Stevenhagen

ECC, Utrecht
September 22, 2008

Abelian varieties and cryptography

They both have a long history – but but their common history is rather short.

1984: Schoof efficiently counts points of elliptic curves over finite fields. Nobody is interested.
(He computed $\sqrt{-1} \mod p$ with it to sell the algorithm.)

1985: Lenstra uses the group of points of an elliptic curve over $\mathbf{Z}/n\mathbf{Z}$ to factor n. Everybody is interested.

Abelian varieties and cryptography

The idea of replacing multiplicative groups by elliptic curves immediately proves to be useful in

- elliptic curve cryptography;
- elliptic curve primality proving.

Complex multiplication naturally enters the scene (ECPP).

Elliptic curves are 1-dimensional abelian varieties.
The extension to higher dimensions is an obvious possibility.

Initially only of theoretical value (Adleman-Huang), but now becoming practical.

What is needed in cryptography?

The discrete logarithm problem (DLP) exists in every group G:

given $x, y \in G$, determine $n \in \mathbf{Z}$ with
$$x^n = y$$
in case such an integer n exists.

In cryptographic protocols such as Diffie-Hellman, n usually exists by construction.
No generality is lost if G is assumed to be abelian or cyclic.

G should be large but finite, with efficient group operations.

Key question: for which G can we guarantee that DLP is 'hard' for most $x, y \in G$?

Generalities on DLP

General algorithms like baby-steps, giant-steps and Pollard-ρ solve DLP in 'arbitrary' G in exponential time, about $\sqrt{\#G}$.

Ideally, we want groups G for which no better algorithms exist.

If we know the group order $\#G$, we can factor it in subexponential time and solve DLP separately in each of the Sylow-p-subgroups of G.

At small p this is easy. We therefore want $\#G$ to be non-smooth, preferably prime or almost prime.

Proving hardness of DLP for concrete G is still out of reach. We are used to working with heuristic run times.

Multiplicative groups

Let **F** be a finite field of order q.

The multiplicative group $\mathbf{F}^*$ is a cyclic group of order $q-1$ that can be used for cryptographic purposes.

Advantage: constructing suitable $\mathbf{F}^*$ is relatively easy.

This is mainly because about one out of every $\log N$ numbers around N is prime by the prime number theorem.

Disadvantage: index calculus provides a subexponential solution to DLP, so q has to be rather large.

Torus based cryptography achieves key sizes reduction by a constant factor.

Groups coming from elliptic curves

Let **F** be a finite field of order q.

The group $E(\mathbf{F})$ of points of an elliptic curve E defined over **F** is of size

$$\#E(\mathbf{F}) \in [(\sqrt{q}-1)^2,\ \ (\sqrt{q}+1)^2]$$

and can be used for cryptographic purposes.

Advantage: no general subexponential solutions to DLP in $E(\mathbf{F})$ are known, so smaller key sizes suffice.

We can costumize E and **F** to meet our demands.
Not all demands can be met so easily...

Constructing elliptic curves

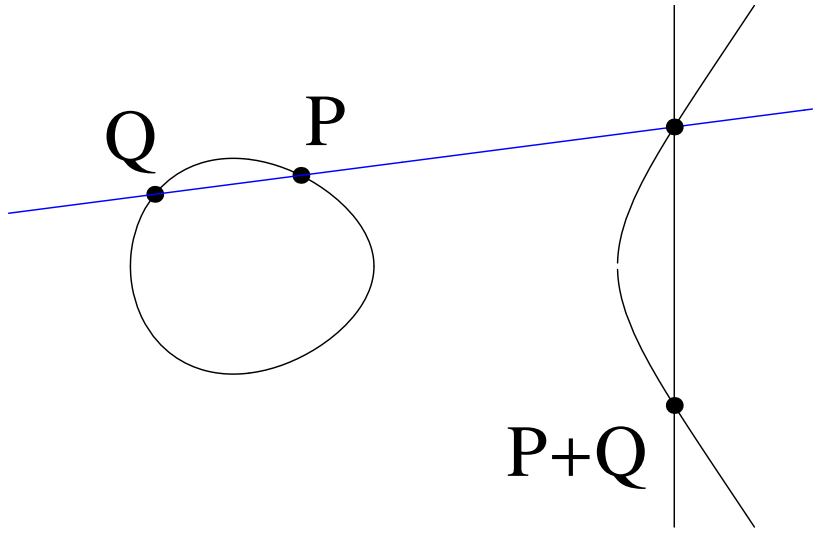

For $p = \mathrm{char}(\mathbf{F}) > 3$, elliptic curves over $\mathbf{F}$ may be given (in $O(\log q)$ bits) by an affine Weierstrass equation

$$Y^2 = X^3 + AX + B \qquad \text{with } A, B \in \mathbf{F} \text{ and } 4A^3 + 27B^2 \in \mathbf{F}^*.$$

The set $E(\mathbf{F})$ of solutions in $\mathbf{P}^2(\mathbf{F})$ naturally forms a group.

The order of $E(\mathbf{F})$

Let us assume for simplicity that $\mathbf{F} = \mathbf{F}_p$ is a prime field.

Determining the order $N = \#E(\mathbf{F})$ efficiently from a Weierstrass equation for E is non-trivial;
this is the point counting done by Schoof's algorithm.

The order N is an integer in the Hasse interval

$$\mathcal{H}_p = [p + 1 - 2\sqrt{p}, p + 1 + 2\sqrt{p}].$$

Conversely, every $N \in \mathcal{H}_p$ arises as the order of some $E/\mathbf{F}$.

The Frobenius endomorphism

The key object that controls the arithmetic properties of an elliptic curve E over $\mathbf{F} = \mathbf{F}_p$ is the Frobenius endomorphism

$$\begin{array}{ccc} E & \longrightarrow & E \\ (X, Y) & \longmapsto & (X^p, Y^p). \end{array}$$

In the endomorphism ring $\mathrm{End}(E)$ of E, it satisfies a quadratic equation

$$\mathrm{Fr}^2 - t \cdot \mathrm{Fr} + p = 0$$

of discriminant $D = t^2 - 4p < 0$.

The Frobenius endomorphism (2)

The ring $\mathbf{Z}[\mathrm{Fr}]$ 'is' an imaginary quadratic order $\mathcal{O}_D$ of discriminant $D = t^2 - 4p$, in which the Frobenius element π satisfies $\pi\overline{\pi} = p$.

$$\begin{aligned} \mathbf{Z}[\mathrm{Fr}] &\xrightarrow{\sim} \mathcal{O}_D = \mathbf{Z}\left[\frac{D+\sqrt{D}}{2}\right] \\ \mathrm{Fr} &\longmapsto \pi = \frac{t+\sqrt{D}}{2}. \end{aligned}$$

If E is ordinary, then $\mathbf{Z}[\mathrm{Fr}]$ is of finite index in $\mathrm{End}(E)$.

Note that D and p determine t up to sign.

(We disregard the supersingular case $t = 0$.)

The trace of Frobenius

Determining $N = \#E(\mathbf{F})$ amounts to computing the trace of Frobenius $t \in \mathbf{Z}$ in the characteristic polynomial

$$f_{\mathbf{Q}}^{\pi} = T^2 - t \cdot T + p$$

of the Frobenius endomorphism as we have

$$N = \#\ker[1 - \mathrm{Fr}] = \mathrm{Norm}(1 - \pi) = p + 1 - t.$$

Schoof's algorithm computes t mod ℓ for many small primes ℓ, and finds t (and N) in polynomial time from E.

Elliptic curve construction

One needs an algorithm in the opposite direction to construct curves $E/\mathbf{F}_p$ for which N (or t) has a prescribed value.

This amounts to finding $E/\mathbf{F}_p$ with complex multiplication by $\mathcal{O}_D$, with $D = t^2 - 4p$. Such E have $\#E(\mathbf{F}_p) = p + 1 \pm t$.

It suffices to find the j-invariant of $E : Y^2 = X^3 + AX + B$, which is defined as $1728 \cdot 4A^3/(4A^3 + 27B^2)$.

Given $j_0 \neq 0, 1728$, the curve

$$E_C : \qquad Y^2 = X^3 + CX - C$$

has j-invariant j_0 for $C = \frac{27j_0}{4(1728-j_0)}$, and $(1,1) \in E_C(\mathbf{F}_p)$.

The j-invariant determines E over $\mathbf{F}_p$ up to quadratic twist.

Complex multiplication

The j-invariants of the complex elliptic curves with endomorphism ring $\mathcal{O}_D \subset \mathbf{C}$ can be computed by complex analytic means.

As Riemann surfaces, they are of the form $\mathbf{C}/\mathfrak{a}$ for an invertible $\mathcal{O}_D$-ideal $\mathfrak{a}$. (Yes, the doughnut...)

Their isomorphism classes correspond to the ideal classes in $\mathrm{Cl}(\mathcal{O}_D)$, which were enumerated by Gauss in terms of binary quadratic forms of discriminant D.

There are about $|D|^{1/2}$ of them.

Complex multiplication (2)

The class polynomial

$$H_D = \prod_{[\mathfrak{a}] \in \mathrm{Cl}(\mathcal{O}_D)} (X - j(\mathfrak{a})) \in \mathbf{Z}[X]$$

has integral coefficients, so it can be computed exactly and may be reduced modulo p.

The polynomial H_D splits into linear factors in $\mathbf{F}_p[X]$.

Its roots in $\mathbf{F}_p$ are the j-invariants of the elliptic curves over $\mathbf{F}_p$ having CM by $\mathcal{O}_D$. Up to twisting, they are all isogenous and have $p + 1 \pm t$ points.

Complex multiplication (3)

Problem:

- H_D has degree $\widetilde{O}(|D|^{1/2})$;
- its coefficients require $\widetilde{O}(|D|^{1/2})$ bits.

It takes time $O(|D|^{1+\varepsilon})$ to compute (and write down) H_D.

Current algorithmic practice: $|D| \lesssim 10^{12}$ (Sutherland).

For **most** values of t, the discriminant $D = t^2 - 4p$ will be as large as p, so the runtime of this CM-method is exponential.

Efficient general curve construction for pairs (p, N) remains a fundamental open problem.

Elliptic curves of prime order

The Schoof-Elkies-Atkin point counting method has become sufficiently efficient to find 'cryptographic curves' of prime order over $\mathbf{F}_p$ by trial and error, in heuristic time $\widetilde{O}((\log p)^5)$.

Theorem (Bröker-S., Contemp. Math. 468 (2008))

On input of a prime number N, one can use the CM-method to construct a finite field $\mathbf{F} = \mathbf{F}_p$ and an elliptic curve E over $\mathbf{F}$ satisfying

$$\#E(\mathbf{F}) = N$$

in heuristic time $\widetilde{O}((\log N)^3)$.

The algorithm is fast enough to handle primes of a few thousand decimal digits.

Sketch of the algorithm

We need to find a quadratic order $\mathcal{O}_D$ with **small** D in which there exists a prime element π for which we have

$$\text{Norm}(1 - \pi) = N.$$

This means that N splits in $\mathcal{O}_D$ as $N = \nu\overline{\nu}$ with

$$\text{Norm}(1 \pm \nu) = p \quad \text{(prime)}.$$

- build up 'small' D from prime discriminants $\pm s \equiv 1 \bmod 4$ that are squares modulo N; store their square roots;
- split $N = \mathfrak{n}\overline{\mathfrak{n}}$ into primes by computing $(\sqrt{D} \bmod N)$;
- test principality of $\mathfrak{n}$ with Cornacchia's algorithm;
- for principal primes $\nu\mathcal{O}_D$, if $p = \text{Norm}(1 \pm \nu)$ is a probable prime, find H_D and (probably) the desired curve.

Heuristic analysis

Heuristic basis:

- numbers Norm$(1 \pm \nu)$ around N will be prime with 'probability' $1/\log N$.
- primes in quadratic orders $\mathcal{O}_D$ will be principal with 'probability' $1/$class number.

Deduce that we will succeed for D of size $\widetilde{O}((\log N)^2)$, and derive the run time.

High level description:
first use the arithmetic in quadratic orders to come up with an appropriate prime element representing Frobenius, then construct an elliptic curve with that Frobenius using CM.

Genus 2 analogues

Much of the theory of elliptic curves has a genus 2 analogue. Smooth projective genus 2 curves (take $\mathrm{char}(k) \neq 2, 3$) look like

$$C : Y^2 = f(X) \in k[X] \qquad \text{with } \deg(f) \in \{5, 6\}.$$

The analogue of the Legendre normal form of elliptic curves is the Rosenhain form

$$Y^2 = X(X-1)(X-\lambda_1)(X-\lambda_2)(X-\lambda_3).$$

It shows that the moduli space of genus 2 curves is 3-dimensional rather than 1-dimensional.

Genus 2 analogues (2)

The isomorphism class (over $\overline{k}$) of a genus 2 curve is determined by the (absolute) Igusa invariants i_1, i_2 and i_3 that are symmetric expressions in the roots of the polynomial f defining C, and lie in k.

Conversely, for every triple $(i_1, i_2, i_3) \in k^3$ of Igusa invariants (with $i_1 \neq 0$) there exists a genus 2 curve C with these invariants.

Computing C from its Igusa invariants is non-trivial (Mestre's algorithm), and C may only be defined over a quadratic extension of k.

The Jacobian

For a genus 2 curve, the k-valued points of C do not naturally form a group.

We do have a group $\mathrm{Jac}(C)$ of divisor classes of degree 0 on C, the Jacobian of C.

Elliptic curves coincide with their Jacobian under the Abel-Jacobi map $P \mapsto [(P) - (\infty)]$.

In genus 2 the map $C \longmapsto \mathrm{Jac}(C)$, embeds the curve C into the abelian surface $\mathrm{Jac}(C)$.

Genus 2 addition law

$$[P_1+P_2-2\infty]+[Q_1+Q_2-2\infty] = -[R_1+R_2-2\infty] = [S_1+S_2-2\infty]$$

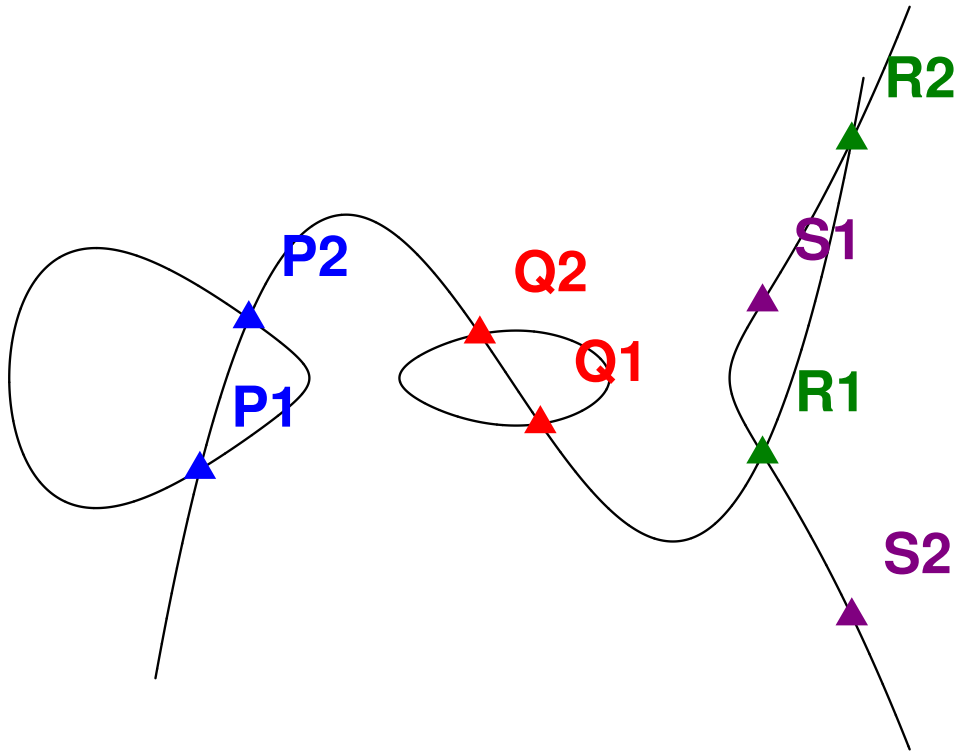

$$\mathrm{div}(y-p(x)) = P_1+P_2+Q_1+Q_2+R_1+R_2-6\infty$$

$$\mathrm{div}(x-x(R_i)) = R_i+S_i-2\infty$$

Hyperelliptic curve cryptography

We can replace the group of points of an elliptic curve E over a finite field $\mathbf{F} = \mathbf{F}_p$ by the group of $\mathbf{F}$-valued points of the Jacobian $J = \mathrm{Jac}(C)$ of a genus 2 curve C.

The order of the group $J(\mathbf{F})$ is an integer in the interval

$$\mathcal{H}_{p,2} = [(\sqrt{p}-1)^4, (\sqrt{p}+1)^4]$$

around p^2.

Can we use $J(\mathbf{F})$ for cryptographic purposes?
Not surprisingly, no general subexponential solution to the DLP in $J(\mathbf{F})$ is known.

The Frobenius

As in the case of elliptic curves, the Frobenius endomorphism controls much of the arithmetic of $J(\mathbf{F})$.

The characteristic polynomial of Frobenius is now of the form

$$f_{\mathbf{Q}}^{\pi} = X^4 + aX^3 + (b+2p)X^2 + apX + p^2 \in \mathbf{Z}[X]$$

for integers a, b satisfying $|a| \leq 4\sqrt{p}$, $|b| \leq 4p$,
and the order of $J(\mathbf{F})$ equals

$$f_{\mathbf{Q}}^{\pi}(1) = \mathrm{Norm}(1-\pi) = (p+1)^2 + a(p+1) + b \in \mathcal{H}_{p,2}.$$

Point counting and CM-method

The Schoof-type algorithm for point counting in genus 2 is more complicated (cf. Schost's talk), but cryptographic size is now getting within reach.

(Schost: 3000 times a month will do...)

This will enable us to perform trial-and-error constructions.

There is also a CM-method to construct genus 2 curves over $\mathbf{F}_p$ with prescribed Frobenius polynomial $f_{\mathbf{Q}}^{\pi}$.
It was pioneered by Weng (2003).

If the quartic CM-field $K = \mathbf{Q}(\pi)$ is small, one can compute the Igusa class polynomials of $\mathcal{O}_K$.

Kohel has an expanding database (Echidna) listing them.

Igusa class polynomials

The Igusa class polynomials of a (primitive) quartic CM-field K are the polynomials

$$H_{K,n} = \prod_C (X - i_n(C)) \in \mathbf{Q}[X] \qquad (n = 1, 2, 3),$$

where C ranges over the complex genus 2 curves for which the endomorphism ring equals $\mathcal{O}_K$.

They are much harder to compute than the class polynomials in genus 1.

Until recently, no run times had been proven.

CM-method in genus 2

Theorem (Streng, preprint on homepage (2008))
The polynomials $H_{K,n}$ have bit size $\widetilde{O}(\Delta_K^2)$ and can be computed from K in time $\widetilde{O}(\Delta_K^{7/2})$.

Here Δ_K is the discriminant of K, and the CM-field K is provided in the form

$$K = \mathbf{Q}(\sqrt{\Delta_0}, \sqrt{-a + b\sqrt{\Delta_0}}) \qquad \text{with } 0 < a, b < \Delta_K.$$

Neither of these bounds is expected to be sharp.

As the degree of $H_{K,n}$ grows like a power of the discriminant, the algorithm is intrinsically exponential.

Very brief sketch of proof

The proof improves upon the work of Spallek (1994), van Wamelen (1999), Weng (2003) and Dupont (2006), and uses the published (2007) und unpublished denominator bounds for Igusa class polynomials of Goren and Lauter.

It computes $H_{K,n}$ from complex approximations of the roots. The proof includes

- computing a list of isomorphism classes of principally polarized abelian varieties $A = \mathbf{C}/\mathfrak{a}$ having CM by $\mathcal{O}_K$;
- computing $d \in \mathbf{Z}_{>0}$ such that $dH_{K,n}$ is in $\mathbf{Z}[X]$;
- computing corresponding period matrices Z in the Siegel upper half space, and moving them under $\mathrm{Sp}_4(\mathbf{Z})$ into (or close to) the fundamental domain;
- computing upper and lower bounds for the theta constants $\vartheta[c](Z)$ needed to compute $i_n(A)$;
- an analysis of the needed precision all along the way.

Abelian surfaces of prime order

If point counting becomes sufficiently fast, we can construct abelian surfaces of prime order that are cryptographically secure by simple trial and error.

This means that we can prescribe the order of magnitude of the desired group order $N = J(\mathbf{F})$, but not N itself.

It may remain faster to use Weng's method in combination with Kohel's database.

Abelian surfaces of prime order (2)

We cannot hope for a theorem as nice as for elliptic curves.

Theorem (Howe-Lauter-S.)

The CM-method does not allow a polynomial time algorithm to construct, on input of a prime number N, a finite field $\mathbf{F} = \mathbf{F}_p$ and an abelian surface J over $\mathbf{F}$ having $\#J(\mathbf{F}) = N$.

The reason is actually simple: there are not enough 'small' quartic CM-fields to deal with all the prime values N below a given bound.

Higher dimensional abelian varieties

Unlike elliptic curves, higher-dimensional abelian varieties are not in general defined by simple equations, and do not possess an explicit algebraic group structure.

Complex analytically, they arise as tori $\mathbf{C}^g/\Lambda$ for $2g$-dimensional lattices Λ that admit a polarization.
They can be embedded as algebraic varieties in high-dimensional projective spaces using theta-functions.

They are not in general Jacobians in dimension ≥ 4.
If they are Jacobians in dimension ≥ 3, they may be Jacobians of non-hyperelliptic curves.

Algorithmically speaking, the CM-method has not been developed beyond dimension $g = 3$.

Weil numbers

It is nevertheless sometimes possible to construct abelian varieties A of higher dimension over finite fields $\mathbf{F} = \mathbf{F}_q$ with 'good' cryptographic properties.

This is because we can conveniently study these in terms of their Frobenius endomorphisms, which are Weil q-numbers $\pi \in \overline{\mathbf{Q}}$.

This means that π has absolute value $\sqrt{q}$ under every complex embedding $\mathbf{Q}(\pi) \to \mathbf{C}$.

Honda-Tate theory

Weil q-numbers (up to conjugation) correspond bijectively to isogeny classes of simple abelian varieties A over $\mathbf{F}_q$.

The correspondence is $\pi \leftrightarrow \mathrm{Fr}_A$.

For a Weil q-number $\pi \neq \pm\sqrt{q}$ the field $\mathbf{Q}(\pi)$ is a CM-field of degree $2g$ with g the dimension of the corresponding abelian variety A.

We have $\#A(\mathbf{F}_q) = \mathrm{Norm}(1-\pi)$.

Pairing-friendly abelian varieties

Weil numbers can be constructed to prove existence of abelian varieties A over $\mathbf{F}_q$ with pleasant properties.

Not only the order $\#A(\mathbf{F}_q)$ can be controlled.

Suppose one fixes (cf. this morning's notation):

- a CM-field K of degree $2g \geq 4$;
- a positive integer $k > 0$;
- a prime $r \equiv 1 \bmod k$ that splits completely in K.

Pairing-friendly abelian varieties

- a CM-field K of degree $2g \geq 4$;
- a positive integer $k > 0$;
- a prime $r \equiv 1 \bmod k$ that splits completely in K.

Theorem (Freeman-S.-Streng, ANTS 2008)

For fixed K, one can find in time polynomial in $\log r$*:*

- *a prime q;*
- *a Weil q-number π such that the corresponding abelian variety $A/\mathbf{F}_q$ has embedding degree k with respect to r.*

The last condition means r divides $\#A(\mathbf{F}_q)$, and that the cyclotomic extension $\mathbf{F}_q \subset \mathbf{F}_q(\zeta_r)$ has degree k.

Pairing-friendly abelian varieties(2)

Basic idea:

- create integers $\pi \in \mathcal{O}_K$ satisfying $\pi\overline{\pi} \in \mathbf{Z}$ by taking

$$\pi = \mathrm{Norm}_{\Phi}(\xi)$$

  for an algebraic integer ξ in the reflex field $\widehat{K}$ of K under the type norm $\mathrm{Norm}_{\Phi} : \widehat{K} \to K$.
- impose congruence conditions on $\xi \in \widehat{K}$ modulo the primes over r to obtain $r \mid \mathrm{Norm}(1-\pi)$ and guarantee that $\pi\overline{\pi} = q$ has order k in $\mathbf{Z}/r\mathbf{Z})^*$.
- take small lifts $\xi \in \mathcal{O}_{\widehat{K}}$ and test whether the resulting number $q = \pi\overline{\pi}$ is prime.

Pairing-friendly abelian varieties(3)

In genus 2, we can combine this with the CM-method to perform actual constructions of pairing friendly genus 2 Jacobians.

In this case, the quotient

$$\rho = \frac{g \log q}{\log r}$$

will lie around 8 without further optimization.
Optimal choices of ξ bring it close to 4.

Wine and cheese?