Faster elliptic-curve discrete logarithms on FPGAs

Daniel J. Bernstein[1,2], Susanne Engels[3], Tanja Lange[2], Ruben Niederhagen[2], Christof Paar[3], Peter Schwabe[4], and Ralf Zimmermann[3]

[1] Department of Computer Science
University of Illinois at Chicago, USA
djb@cr.yp.to
[2] Department of Mathematics and Computer Science
Eindhoven University of Technology, The Netherlands
tanja@hyperelliptic.org, ruben@polycephaly.org
[3] Horst Görtz Institute for IT-Security
Ruhr-University Bochum, Germany
susanne.engels@rub.de, christof.paar@rub.de, ralf.zimmermann@rub.de
[4] Institute for Computing and Information Sciences
Radboud University, The Netherlands
peter@cryptojedi.org

Abstract. This paper accelerates FPGA computations of discrete logarithms on elliptic curves over binary fields. As a toy example, this paper successfully attacks the SECG standard curve sect113r2, a binary elliptic curve that was not removed from the SECG standard until 2010 and was not disabled in OpenSSL until June 2015. This is a new size record for completed ECDL computations, using a prime order very slightly larger than the previous record holder. More importantly, this paper uses FPGAs much more efficiently, saving a factor close to 3/2 in the size of each high-speed ECDL core. This paper squeezes 3 cores into a low-cost Spartan-6 FPGA and many more cores into larger FPGAs. The paper also benchmarks many smaller-size attacks to demonstrate reliability of the estimates, and covers a much larger curve over a 127-bit field to demonstrate scalability.

Keywords: attacks, FPGAs, ECC, binary curves, Pollard rho, negation

1 Introduction

FPGAs are the most energy-efficient mass-market devices for computations of discrete logarithms on elliptic curves defined over binary fields. For example, in

Public domain. This work was supported by NSF (U.S.) under grant 1018836; by NWO (Netherlands) under grants 639.073.005, 613.001.011, and Veni 2013 project 13114; and by the European Commission through the ICT program under contract INFSO-ICT-284833 (PUFFIN). Permanent ID of this document: 01ac92080664fb3a778a430e028e55c8. Date: 2016.08.06.

2009 a large cross-platform effort was initiated [1] to attack Certicom's ECC2K-130 challenge; an implementation optimized for a 300-watt NVIDIA GTX 295 GPU (dual 55nm GT200) performed 63 million iterations per second [3], while an implementation of the same iteration function optimized for a 5-watt Xilinx XC3S5000 (90nm Spartan-3) FPGA performed 111 million iterations per second [15], obviously much better performance.

An FPGA is a well-connected mesh of a large number of programmable "lookup tables" (LUTs) surrounded by other useful resources such as "registers". For example, an XC3S5000 contains 74880 "LUT-4" units; each LUT-4 maps 4 input bits to 1 output bit. This paper focuses on newer Xilinx FPGAs with larger "LUT-6" units, each mapping 6 input bits to 1 output bit: e.g.,

- a Xilinx XC6SLX150 (45nm Spartan-6, typically clocked at 100MHz) contains 92152 LUT-6; and
- a Xilinx XC7K325T-2 (28nm Kintex-7, typically clocked at 180MHz) contains 203800 LUT-6.

A modern high-end GPU is also highly parallel, with thousands of 32-bit arithmetic units, but most of the operations provided by these units (e.g., floating-point multiplication) are not helpful for binary-field arithmetic, and the remaining operations (e.g., xor) are quite wasteful, spending most of their energy and chip area on data transfer rather than computation. The LUTs in FPGAs handle binary-field computations with much less overhead.

Recent binary-field elliptic-curve discrete-log work by Wenger and Wolfger (see the preliminary SAC 2014 paper [31] and the final journal paper [32]) fits 5 cores onto an XC7K325T-2, using a total of 151 KLUTs, i.e., 30 KLUTs per core. Each core runs at 180MHz and computes 1 iteration per cycle, so in total the FPGA computes 900 million iterations per second. Wenger and Wolfger used a cluster of ten KC705 development boards, each with an XC7K325T-2 FPGA, for 2.5 months to successfully compute a discrete logarithm on an elliptic curve defined over the field $\mathbb{F}_{2^{113}}$. The prime order here was slightly above 2^{112}, not much larger than the prime order from the previous ECDL record (approximately $2^{111.78}$; see [7] and [8]), but the previous ECDL record used PlayStations and would have needed about 300 PlayStations to be completed in 2.5 months.

Wenger and Wolfger refer to several prior ECDL implementations on FPGAs but claim in [32] that "none of their FPGA implementations have been successful in solving ECDLPs". This is contradicted by, e.g., [14, Section 5.4], which reported "successfully" breaking a "target with a 60-bit ECDLP"; note that [14] was cited as [32, reference 14]. It was already clear that FPGAs are very efficient for this task; the remaining question is whether they can be sped up even more.

1.1. Primary contribution of this paper: more efficient ECDL cores. In this paper we do better than [32] by a factor around 3/2: we squeeze ECDL cores into just 21 KLUTs per core for the same field $\mathbb{F}_{2^{113}}$, while maintaining high clock frequencies and maintaining a speed of 1 iteration per cycle on each core. For example, we fit 6 cores into 126 KLUTs, 7 cores into 145 KLUTs, and 8 cores into 163 KLUTs. This speedup combines three directions of improvements:

- Smaller high-speed multipliers. Our $\mathbb{F}_{2^{113}}$ multiplier takes just 3071 LUTs. The multiplier in [32] takes 3757 LUTs, 22% larger.
- Fewer multipliers. For example, we use 16 multipliers for 3 cores and 32 multipliers for 6 cores, while the approach of [32] needs 15 multipliers for just 2 cores and 30 multipliers for just 5 cores.
- Reduced area outside the multipliers. For example, we are the first to point out that $\mathrm{Tr}(x) = 1$ allows a new lower-area definition of $|P|$. The total number of LUTs we use for the entire iteration function is only about 30% more than the number used for the multiplications. For [32] the overhead is around 50%.

We do not have access to any Kintex-7 FPGAs for testing, but we do have access to low-cost Spartan-6 FPGAs. We fit 3 cores into just 64 KLUTs and tested those cores at 100MHz on an XC6SLX150; this is a total of 300 million iterations per second, achieving 1/3 the speed of [32] using an FPGA that costs only 1/5 as much.

We estimate that scaling the design of [32] down to 1 core would make it fit into an XC6SLX150 and would successfully run at 100MHz, but this would compute only 100 million iterations per second. It is not at all clear that 2 cores would fit: 15 multipliers at 3757 LUTs would already consume 56355 LUTs even without counting overhead. It is clear that 3 cores would not fit. For our design 3 cores fit easily. On the more powerful Kintex-7 used in [32], we expect our 7-core design to run stably at 180MHz, computing 1.26 billion iterations per second on an XC7K325T-2, 40% faster than [32]. Our 8-core design also fits, and if it runs stably at 180MHz then it is 60% faster than [32]. See Section 6.3 for a more detailed stability analysis.

To put this 3/2 speedup into perspective, note that ECDL has for many years been viewed as a highly optimized cryptanalytic computation, with very little room for improvement. The number of iterations to compute an average ECDL on a generic curve has for many years been asymptotically $(1 + o(1))\sqrt{\pi\ell/4}$, where ℓ is the prime order. The last major improvement here was the negation map, which has been the topic of several ECDL papers and saved only $\sqrt{2} - o(1) < 1.5$. The amount of arithmetic per iteration has for many years been asymptotically the cost of $5 + o(1)$ binary-polynomial multiplications.[5] The remaining questions are how much the $o(1)$ overheads can be reduced and how efficiently the multiplications can be carried out in a given amount of hardware.

See Sections 3 and 4 for further details of our design, and Section 6 for analysis of how our results improve upon the results of [32].

[5] Each inversion costs $3 + o(1)$ multiplications by Montgomery's trick, so each division costs $4 + o(1)$ multiplications. The resulting λ is used in 1 multiplication. The squaring of λ costs asymptotically only as much as $o(1)$ multiplications since we are using a binary field, and similar comments apply to additions, reductions, comparisons, etc., for a total of $5 + o(1)$ multiplications.

We point out that this well-known $5 + o(1)$ can be improved to $4.5 + o(1)$ as follows. Choose the ith table entry R_i to have x-coordinate matching the bits of i at the positions that are used to select i. This forces each denominator $x_1 + x_2$ used in λ to have bits 0 in all of those positions. Use a precomputed table of maximum size

1.2. Secondary contribution of this paper: a new ECDL record. The curve sect113r2 was standardized by SECG (Standards for Efficient Cryptography Group) in version 1.0 of "SEC 2: Recommended Elliptic Curve Domain Parameters" [10] in 2000, and was included as one of the supported curves when OpenSSL added ECC support in 2005 (version 0.9.8). This curve is defined over $\mathbb{F}_{2^{113}}$; see Section 2 for further details of the curve.

This curve disappeared from version 2.0 of the SEC 2 standard [11] in 2010 (along with all other curves over field sizes below $\approx 2^{192}$ for odd characteristic and 2^{163} for even characteristic), and disappeared from OpenSSL's default curve list in June 2015 (version 1.0.2b). However, most installations are running older versions of OpenSSL that still support this curve. It is easy to imagine how an OpenSSL user seeking to minimize bandwidth for ECC would look through the output of openssl ecparam -list_curves and find this curve as one of the lowest-bandwidth options.

We used 120 XC6SLX150 FPGAs to successfully compute an ECDL of a random target point on sect113r2. This computation was slower than desired for three reasons: first, it used a preliminary 2-core version of our implementation; second, not all of the FPGAs were available all of the time; third, the number of iterations in these ECDL algorithms is a random variable with high variance, and we were moderately unlucky in the number of iterations used for this particular computation. The computation nevertheless finished in under 2 months.

Technically, this is a new ECDL record, since the prime order is (marginally) larger than the prime order in [32]. We do not mean to exaggerate the importance of setting ECDL size records; obviously such records are heavily influenced by hardware availability, obscuring the impact of algorithmic improvements and understating the amount of hardware actually available to attackers. We also do not mean to exaggerate the importance of this curve: the curve is toy crypto that should never have been standardized. We do not know how many people have actually used sect113r2 to encrypt data. What really matters in this paper is being able to squeeze iterations into fewer LUTs (see Section 1.1), reducing costs not merely for this attack but also for much larger attacks against much larger curves.

1.3. Variations and extrapolations. The extrapolations in [32, Section 7.1 and Table 4] assume that the number of FPGA-years scales as a simple square root of the prime order. This assumption means, for example, that the NIST B-163 prime order (almost exactly 2^{162}) costs 2^{25} times as many FPGA-years as the sect113r1 and sect113r2 prime orders (almost exactly 2^{112}). The extrapolations also assume a $\sqrt{163}$ speedup for the Koblitz curve NIST K-163.

However, an accurate cost analysis is more complicated. Many components of an ECDL core grow linearly with the number of bits in the field. Even worse, the

$\ell^{0.5-o(1)}$, so that each denominator has only $(0.5+o(1))\log_2 \ell$ bits. This reduces the cost of denominator multiplication by a factor $2+o(1)$.

Unfortunately, this asymptotic improvement in the amount of arithmetic costs too much area to be useful in a more sophisticated cost metric. See later for details of our table usage and of how we merge point doublings with point additions.

area for a high-speed multiplier grows superlinearly. There is also a noticeable extra reduction cost for fields defined by pentanomials rather than trinomials, such as $\mathbb{F}_{2^{163}}$. Scaling to larger and larger fields will eventually force any particular size of FPGA to use fewer and fewer high-speed cores.

We scaled our design from a 113-bit field up to a 127-bit field, carefully reoptimizing our Karatsuba-based arithmetic. This expanded 3 cores from 64401 LUTs to 74095 LUTs. We tested that this 3-core design still fits onto an XC6SLX150 and runs successfully at 100MHz. We are now running this design on 128 of these FPGAs to compute an ECDL on a curve over $\mathbb{F}_{2^{127}}$ in a subgroup of prime order approximately $2^{117.35}$; this is expected to take only 123 days. For comparison, as noted above, it is not at all clear that 2 of the 113-bit cores from [32] would fit onto this FPGA, and it seems quite unlikely that 2 similar 127-bit cores would fit.

The bigger picture is that attacks at interesting sizes should be less expensive than predicted in [32]. The caveat that cores grow with field size is outweighed by our 3/2 improvement considerably beyond 127 bits. Furthermore, given the agility of FPGAs to promptly and cost-effectively tackle new problems, any serious attacker should be expected to be operating a large FPGA cluster; and, given economies of scale, the cost per FPGA in a large cluster should be expected to be much lower than indicated in [32, Table 4]. We do not agree, for example, that a 5-million-FPGA cluster for breaking NIST K-163 in a year would cost $10 \cdot 10^9$ USD, an entire year of NSA's budget. A more plausible estimate is under $2 \cdot 10^9$ USD, similar in cost (and power consumption) to one of NSA's existing data centers.

A well-funded attacker facing years of predictable large-scale computations will do even better by building application-specific integrated circuits (ASICs). FPGA optimization techniques are well known to be much better than CPU (and GPU) optimization techniques as a predictor of ASIC optimization techniques. Our multiplier details should be reoptimized for ASICs but we expect the overall architecture to perform very well.

1.4. Attacking many targets. We use Q-independent walks (see Section 3), so distinguished points collected in solving one ECDL help solve the next ECDL more quickly. It is well known that this trick breaks K keys at cost only about $\sqrt{K}$ times as much as breaking one key; see [21], [20], and [4].

For example, a cluster breaking K-163 in one year would be expected to break approximately 25 keys, not just 5 keys, in 5 years. This makes a large ECDL cluster more attractive for the attacker, and more damaging for the users. Furthermore, our 3/2 speedup in finding distinguished points (for the same hardware cost) actually means that we can break *more than twice as many keys* in the same amount of time.

1.5. Binary fields vs. prime fields. The standard NIST curves fall into three different categories: Koblitz curves over binary fields $\mathbb{F}_{2^n}$, "random" curves over binary fields $\mathbb{F}_{2^n}$, and "random" curves over prime fields $\mathbb{F}_p$. There are five NIST curves in each category, spread across five different sizes of 2^n or p. The smallest 2^n is 2^{163}, while the smallest p is approximately 2^{192}. See [22] and [23].

It is easy to see how an ASIC designer concerned with the costs of an ECC coprocessor (chip area, power, energy, etc.) for constructive use will end up choosing $\mathbb{F}_{2^{163}}$. Taking a binary field rather than a prime field eliminates all the circuitry for carries, and taking the smallest allowable field has obvious performance benefits. It is not as clear whether the designer will prefer a Koblitz curve or a "random" curve: the extra endomorphisms in Koblitz curves reduce the number of field multiplications inside scalar multiplication, saving energy, but managing these endomorphisms comes at a cost in chip area.

There are several common arguments that prime fields should be preferred, but it is also easy to imagine counterarguments from the ASIC designer:

- Prime fields provide better software performance, since they take better advantage of the integer multipliers provided by CPUs. Counterarguments: CPUs evolve to meet the needs of applications, and Intel's new PCLMULQDQ instruction already provides excellent performance for curves over binary fields; see [25]. Software performance is ultimately less important than hardware performance.
- Prime fields are the safest choice, since binary fields have extra structure that might be exploitable. Some papers have suggested the possibility of an asymptotically subexponential ECDL algorithm for curves over $\mathbb{F}_{2^n}$. Counterarguments: Other papers have disputed this possibility; see [17, Section 10.2] for a recent overview and references. None of the papers have claimed relevance to the range of n actually used in ECC. Speculations about security problems are less important than meeting the performance requirements of the applications.
- For Koblitz curves there are extra endomorphisms that speed up known attacks by a factor close to $\sqrt{n}$. See [33] and [18]. Counterarguments: This speedup does not apply to "random" curves, and $\sqrt{n}$ is a limited factor in any case.
- Some ECC standards require prime fields: consider, e.g., NSA's Suite B [24] and the Brainpool standard [13], both from 2005. Counterargument: Those standards do not articulate reasons to avoid binary fields: e.g., [13] says that subsequent editions "may also contain elliptic curves over fields of characteristic 2".

In this paper we do not take a position in this debate. We merely observe that the performance of binary-field ECC continues to attract attention, so the community also needs to understand the cost of solving binary-field ECDLP. Standard extrapolations (see above) suggest that breaking a "random" curve over $\mathbb{F}_{2^{163}}$ is an order of magnitude more expensive than breaking a Koblitz curve over $\mathbb{F}_{2^{163}}$, which in turn is millions of times more expensive than breaking a random curve over $\mathbb{F}_{2^{113}}$; but these are not infeasible computations, and a 3/2 speedup has a huge impact at this scale. Some of our area-optimization techniques are also applicable to prime-field ECDLP, although obviously the details of arithmetic will be different.

1.6. Acknowledgements. The authors would like to thank Bo-Yin Yang of Academia Sinica for providing access to his FPGA clusters for all computations and to SciEngines for doing computations on their clusters for the 117-bit DLP.

2 Two target curves: sect113r2 and target117

This section describes two sample curves where we are performing computations. The first curve, sect113r2, is from the SECG standard. The second curve, target117, is a larger non-standard curve that we define here. This section also includes trace calculations that we exploit later.

Our successful ECDL computation on sect113r2 means that both of the curves standardized by SECG over $\mathbb{F}_{2^{113}}$ have been broken. The next SECG binary field is $\mathbb{F}_{2^{131}}$, skipping $\mathbb{F}_{2^{127}}$. We decided to create target117 as an intermediate target over $\mathbb{F}_{2^{127}}$.

2.1. Arithmetic on binary elliptic curves. For efficiency, curves over binary fields $\mathbb{F}_{2^n}$ are usually chosen to be of the form $y^2 + xy = x^3 + x^2 + b$. Addition of two points (x_1, y_1) and (x_2, y_2) on this curve produces a result (x_3, y_3) with

$$(x_3, y_3) = (\lambda^2 + \lambda + 1 + x_1 + x_2, \quad \lambda(x_1 + x_3) + y_1 + x_3), \text{ where}$$
$$\lambda = \begin{cases} (x_1^2 + y_1)/x_1 & \text{if } P_1 = P_2 \neq -P_2 \\ (y_1 + y_2)/(x_1 + x_2) & \text{if } P_1 \neq \pm P_2 \end{cases}.$$

The negative of a point is $-(x_1, y_1) = (x_1, y_1 + x_1)$ and $(x_1, y_1) + (x_1, y_1 + x_1) = \infty$.

All curves of the form $y^2 + xy = x^3 + x^2 + b$ have a co-factor of 2, with $(0, \sqrt{b})$ being a point of order 2. Varying b varies the group order but the term x^2 means that there is no point of order 4. Essentially all integer orders within the Hasse interval $[2^n + 1 - 2 \cdot 2^{n/2}, 2^n + 1 + 2 \cdot 2^{n/2}]$ that are congruent to 2 modulo 4 are attainable by changing b within $\mathbb{F}_{2^n}$. We use this to generate further elliptic curves with points of medium prime order for testing purposes.

2.2. The first target curve: sect113r2. The SECG curve sect113r2 is defined over $\mathbb{F}_{2^{113}} \cong \mathbb{F}_2[w]/(w^{113} + w^9 + 1)$ by an equation of the form $E : y^2 + xy = x^3 + ax^2 + b$ and basepoint $P = (x_P, y_P)$, where

$$a = \texttt{0x0689918DBEC7E5A0DD6DFC0AA55C7},$$
$$b = \texttt{0x095E9A9EC9B297BD4BF36E059184F},$$
$$x_P = \texttt{0x1A57A6A7B26CA5EF52FCDB8164797}, \text{ and}$$
$$y_P = \texttt{0x0B3ADC94ED1FE674C06E695BABA1D},$$

using hexadecimal representation for elements of $\mathbb{F}_{2^{113}}$, i.e., taking the coefficients in the binary representation of the integer as coefficients of the powers of w, with the least significant bit corresponding to the power of w^0. The order of P is 5192296858534827702972497909952403, which is prime. The order of the curve $|E(\mathbb{F}_{2^{113}})|$ is twice as large.

It is possible to transform the elliptic curve to isomorphic ones by maps of the form $x' = c^2x + u, y' = c^3y + dx + v$. These maps do not change the general shape of the curve (the highest terms are still y^2, x^3, and xy) but allow mapping to the more efficient representation given above. The security among isomorphic curves is identical: the DLP can be transformed using the same equations. Curve arithmetic depends on the value of a and for fields of odd extension degree it is always possible to find an isomorphic curve with $a \in \{0, 1\}$. It is unclear why this optimization was not applied in SECG but we will use it in the cryptanalysis.

For sect113r2 we have $\mathrm{Tr}(a) = 1$ so there is an element $t \in \mathbb{F}_{2^{113}}$ satisfying $t^2 + t + a + 1 = 0$. Now $(x_P, y_P + tx_P)$ is on $y^2 + xy = x^3 + x^2 + b$ for every (x_P, y_P) on E because

$$\begin{aligned}(y_P + tx_P)^2 + x_P(y_P + tx_P) &= y_P^2 + x_Py_P + (t^2x_P^2 + tx_P^2)\\ &= x_P^3 + ax_P^2 + b + (t^2 + t)x_P^2 = x_P^3 + x_P^2 + b.\end{aligned}$$

The specific value for t is given in Appendix A. The base point is transformed to (x_P, y'_P) with

$$y'_P = \texttt{0x17D5618CD2EE81F84FAB74B1EB19F}.$$

2.3. The second target curve: target117. This curve is defined over $\mathbb{F}_{2^{127}} \cong \mathbb{F}_2[w]/(w^{127} + w + 1)$ by an equation of the form $E : y^2 + xy = x^3 + ax^2 + b$ and base point $P = (x_P, y_P)$, where

$$\begin{aligned}a &= \texttt{0x00000000000000000000000000000001},\\ b &= \texttt{0x000000000000000000000000000001AB},\\ x_P &= \texttt{0x3CF9CCD146B5E7440E9632F5D2B49679}, \text{ and}\\ y_P &= \texttt{0x43ED94FD97454C8197B6207C9A23C67E}.\end{aligned}$$

The order of P is $212146114040485326348618959071598183 \approx 2^{117.35}$ which is prime. The order of the curve $|E(\mathbb{F}_{2^{127}})|$ is 802 times larger.

2.4. Trace calculations. Finite fields of characteristic 2 are usually defined using an irreducible polynomial $f \in \mathbb{F}_2[w]$. For our fields, $f - 1$ is an odd polynomial (i.e., $f - 1 = wg(w^2)$ for some polynomial g), and we can prove some properties about the trace of elements.

Theorem 2.1. *Let n be an odd positive integer. Let $f_1, f_3, \ldots, f_{n-2}$ be elements of $\mathbb{F}_2$. Define $f = 1 + f_1w + f_3w^3 + \cdots + f_{n-2}w^{n-2} + w^n \in \mathbb{F}_2[w]$. Assume that f is irreducible. Define α as the image of w in the finite field $\mathbb{F}_2[w]/(f)$. Then $\mathrm{Tr}(\alpha^i) = 0$ for $1 \le i < n$.*

Proof. We start with Newton's identities expressed as the concise equation

$$\frac{f'}{f} = \sum_{i \ge 0} \mathrm{Tr}(\alpha^i)\epsilon^{i+1}$$

in the field $\mathbb{F}_2((\epsilon))$ of Laurent series, where f' is the derivative of f and $\epsilon = 1/w$. For a proof see, e.g., [2].

We will show explicitly that $f'/f \in \epsilon - \epsilon^{n+1} + O(\epsilon^{n+2})$, where $O(\epsilon^k)$ means the set of series of the form $s_k\epsilon^k + s_{k+1}\epsilon^{k+1} + \cdots$. We then simply read off $\mathrm{Tr}(\alpha), \mathrm{Tr}(\alpha^2), \ldots, \mathrm{Tr}(\alpha^{n-1})$ as the coefficients of $\epsilon^2, \epsilon^3, \ldots, \epsilon^n$ respectively in f'/f: all of these coefficients are 0 as claimed.

The hypothesis $f = 1 + f_1w + f_3w^3 + \cdots + f_{n-2}w^{n-2} + w^n$ implies $f' = f_1 + f_3w^2 + \cdots + f_{n-2}w^{n-3} + w^{n-1} = \epsilon(f-1)$, i.e., $f'/f = \epsilon(1-1/f)$. We will show that $1/f \in \epsilon^n + O(\epsilon^{n+1})$, so $1-1/f \in 1-\epsilon^n+O(\epsilon^{n+1})$, so $f'/f \in \epsilon-\epsilon^{n+1}+O(\epsilon^{n+2})$ as claimed.

The same hypothesis implies $\epsilon^n f = \epsilon^n + f_1\epsilon^{n-1} + f_3\epsilon^{n-3} + \cdots + f_{n-2}\epsilon^2 + 1 \in 1 + O(\epsilon)$. Anything in $1 + O(\epsilon)$ also has reciprocal in $1 + O(\epsilon)$: one standard proof uses the Taylor series for $1/(1 + z)$, and another uses the ϵ valuation. Hence $1/(\epsilon^n f) \in 1 + O(\epsilon)$, i.e., $1/f \in \epsilon^n + O(\epsilon^{n+1})$ as claimed. □

If n is odd and x^n+x^m+1 is irreducible then $x^n+x^{n-m}+1$ is also irreducible. One of these two trinomials f meets the requirement of Theorem 2.1 that $f - 1$ be odd. Not every n has an irreducible trinomial, but it is generally believed that each $n \geq 4$ has an irreducible pentanomial, and it seems that for each odd $n \geq 7$ there is an irreducible degree-n pentanomial f such that $f - 1$ is odd.

2.5. Traces of x-coordinates. Cryptographic applications work in a subgroup of prime order. Because this order ℓ is odd, 2 is invertible modulo ℓ, so there exists an s with $2s \equiv 1 \bmod \ell$. This means that each point R in this subgroup of prime order is the double of sR, because $2sR = R$. Seroussi showed in [27] that on $y^2 + xy = x^3 + ax^2 + b$ for any point (x, y) that is the double of any point it holds that $\mathrm{Tr}(x) = \mathrm{Tr}(a)$.

For both fields $\mathbb{F}_{2^n}$ considered here, i.e. $\mathbb{F}_{2^{113}} \cong \mathbb{F}_2[w]/(w^{113} + w^9 + 1)$ and $\mathbb{F}_{2^{127}} \cong \mathbb{F}_2[w]/(w^{127} + w + 1)$, we have just shown that $\mathrm{Tr}(w^i) = 0$ for $1 \leq i < n$ and, of course, $\mathrm{Tr}(1) = 1$. If $x = \sum_{i=0}^{n-1} x_i w^i$ then $\mathrm{Tr}(x) = \sum_{i=0}^{n-1} x_i \mathrm{Tr}(w^i) = x_0$ since the trace is additive. This implies that for our curves having $a = 1$, each point in the subgroup of order ℓ has $x_0 = \mathrm{Tr}(x) = \mathrm{Tr}(a) = 1$, i.e., the least significant bit in the representation of x is 1.

3 Pollard iterations

Our attack uses the parallel version of Pollard's rho algorithm [26] by van Oorschot and Wiener [30] to compute the discrete logarithm of Q to the base P. This algorithm works in a client-server approach.

Each client (in our case each FPGA process) receives as input a point R_0 which is a known linear combination of P and Q, i.e., $R_0 = a_0P + b_0Q$. From this input point it starts a pseudorandom walk, where each step depends only on the coordinates of the current point R_i and preserves knowledge of coefficients a_i, b_i such that $R_i = a_iP + b_iQ$. The walk ends when it reaches a so-called "distinguished point" R_d, where the property of being distinguished is a property

of the coordinates of the point. This distinguished point is then reported to a server together with information that allows the server to obtain a_d and b_d.

The server searches through incoming points until it finds a collision, i.e., two walks that ended up in distinguished points $R_{d_1} = a_{d_1}P + b_{d_1}Q$ and $R_{d_2} = a_{d_2}P + b_{d_2}Q$ that are equal. With very high probability, the coefficients b_{d_1} and b_{d_2} are distinct modulo ℓ, so we can compute the discrete logarithm as $\log_P Q \equiv (a_{d_1} - a_{d_2})/(b_{d_2} - b_{d_1}) \bmod \ell$.

In the following, we describe the construction of our iteration function. We start with a simple version, which does not make use of the negation map, and then modify this walk to perform iterations modulo negation.

3.1. Non-negating walk. Our iteration function follows the standard approach of an additive walk (see e.g. [29]) with some improvements following [6]. We precompute a table $(T_0, \ldots, T_{n-1})$ of random multiples of the base point P; our implementation uses $n = 1024$. Older descriptions often define steps to be combinations of P and Q, but Q is a multiple of P itself, so taking random multiples of P has the same effect and makes the step function independent of the target discrete logarithm. This means the design including the precomputed points can be synthesized for the FPGA and then used to break multiple discrete logarithms. We use a random multiple of the target point Q for the starting point R_0 of a random walk. Our iteration function f is defined as

$$R_{i+1} = f(R_i) = R_i + T_{I(R_i)},$$

where $I(R_i)$ takes the coefficients of $w^{10}, w^9, \ldots, w^1$ of the x-coordinate of R_i, interpreted as an integer. We ignored the coefficient of w^0 because it is 1 for all points (see Section 2) and chose the next 10 least significant bits in order to avoid overlap with the distinguished-point property defined in the next paragraph.

After each iteration, we check whether we have reached a distinguished point. We call a point *distinguished* when the 30 most significant bits of the x-coordinate are zero. If the point is a distinguished point, it is output, otherwise the iteration proceeds.

In the literature, there are two different approaches of how to continue after a distinguished point has been found. The traditional approach is to report the point and the linear combination leading to it and then to simply continue with the random walk. This approach has been used, for example, in [19], [7], [8], [31], and [32]. The disadvantage of this approach is that the iteration function needs to update the coefficients of the linear combination of P and Q (at least the coefficient of P, and also the coefficient of Q with older definitions of steps); in our case this would mean that the FPGAs not only have to perform arithmetic in $\mathbb{F}_{2^{113}}$ but also big-integer arithmetic modulo the 113-bit group order ℓ.

A more efficient approach was suggested in [1] and [6]. Once a distinguished point has been found the walk stops and reports the point. The processor then starts with a fresh input point. This means that all walks have about the same length, in this case about 2^{30} steps. The walks do not compute the counters for the multiples of P and Q; instead they remember the initial multiple of Q in the form of a seed. The server stores this seed and the resulting distinguished

point. Once a collision between distinguished points has been found, we simply recompute the two colliding walks and this time compute the multiples of P. We wrote a non-optimized software implementation based on NTL for this task, which took time on the scale of an hour to recompute the length-2^{30} walks and solve the DLP on sect113r2.

3.2. Walks modulo negation. Like [32] and various earlier papers, we reduce the expected number of iterations for an ECDL computation by computing iterations modulo the efficiently computable negation map. This improvement halves the search space of Pollard rho and thus gives a theoretic speedup of $\sqrt{2}$. The use of the negation map has been an issue of debate: see [9] for arguments against and [6] for an implementation that achieves essentially the predicted speedup.

Changing the walk to work modulo the negation map requires two changes. First, we have to map $\{P, -P\}$ to a well-defined representative. We denote this representative $|P|$. Recall that the least significant bit of x is always 1 (see the trace discussion at the end of Section 2) and that $-(x, y) = (x, x + y)$. We pick the point that has the least significant bit of y being 0 as representative. After each step of the iteration function we inspect the y-coordinate of the reached point R_i and continue with $-R_i$ in case the least significant bit of y is 1. This requires one bit comparison and one field addition.

The traditional approach would instead be to take whichever of R_i and $-R_i$ has a lexicographically smaller y-coordinate. Our approach, relying on the $x_0 = 1$ observation, replaces a lexicographic comparison with a single bit comparison, noticeably reducing area overhead.

Second, we need a mechanism to escape so-called *fruitless cycles*. These mini-cycles stem from the combination of additive walks and walks defined modulo negation. The most basic and most frequent case of a fruitless cycle is a 2-cycle. Such a cycle occurs whenever $I(R_i) = I(R_{i+1})$ and $R_{i+1} = |(R_i + T_{I(R_i)})| = -(R_i + T_{I(R_i)})$. In this case, R_{i+2} is again R_i and the walk is caught in a cycle consisting of R_i and R_{i+1}. The probability of this to occur is $1/(2n)$, where n is the number of precomputed points. There also exist larger fruitless cycles of lengths $4, 6, 8$ etc., but the frequency of those is much lower. See Appendix B.

Bernstein, Lange and Schwabe suggest in [6] detecting fruitless cycles by checking frequently for cycles of length 2 and increasingly less frequently for cycles of higher length. However, they are using a vectorized software implementation where frequent checks for cycles are expensive. We are using an unrolled hardware design; checking for cycles has no impact on the computational throughput and only a small impact on area demand. Also, cycle checking can be done individually for each independent walk in the pipeline without impact on the other walks (in contrast to a vectorized implementation where the same operation must be applied jointly to all walks in the data vectors).

We use a simple 4-bit counter, allowing us to detect all cycles of length up to 16. This prevents practically all infinite loops during the computation: fruitless cycles of length 10, 12, 14 have probabilities approximately $2^{-47.1}$, $2^{-54.5}$, $2^{-61.8}$ respectively (see Appendix B), and fruitless cycles of length $\geq$16 are extremely unlikely. Once a cycle is detected, a deterministic way of leaving the

cycle is required regardless of where the cycle was entered; two independent walks that enter the same cycle at a different entry point must leave the cycle at the same point in order to eventually end in the same distinguished point. Therefore, we record the current minimum x-coordinate in the 16-step cycle-detection window. Whenever we reach a point with a smaller x-coordinate than the current minimum, the cycle counter is reset and the minimum x-coordinate is updated. When we reach the same point, i.e. the same x-coordinate as the stored minimum (given that we are using the negation map we do not need to compare the y-coordinate), then we are in a cycle and have to escape the cycle by doubling the current point. If the counter has an overflow to 0, i.e., we did not encounter a cycle in the last 16 steps, the current minimum is reset to the current x-coordinate and the cycle-detection is restarted.

We use the same criterion for a distinguished point (30 zeros) and the same table of precomputed steps as described in the previous subsection.

3.3. Justification of distinguished-point property. For the sect113r2 curve the expected number of group operations is roughly 2^{56}. Each walk takes about 2^{30} steps to reach a distinguished point and so we expect about 2^{26} distinguished points before we find a collision. This amount of data poses no problem for the host PC and for the I/O part of the hardware. The same criterion is also a good choice for the target117 curve which then requires about 2^{28} distinguished points. For even larger DL computations a less frequent property needs to be chosen. A benefit of relatively short walks is that they are easily recomputed on a PC, which we use for finding the DL after a collision of distinguished points occurs. This also helped in verifying that the FPGA code computed the same walks as a software implementation.

4 Implementation

The main core of the iteration function is a point addition, either the addition of the current state point with a point from the precomputed table or in case a cycle was detected the doubling of the current state point.

Doubling of a point is quite similar to addition of two distinct points (see Section 2 for the standard addition formula) but removes one finite-field addition and includes one extra finite-field squaring. This can easily be expressed using conditional assignments. Figure 4.1 shows Sage code for the iteration function; the point addition/doubling part is in lines 3 to 16. The code doubles the current state point (x, y) in case the double flag is true or adds a point $T = (T_x, T_y)$ from a precomputed table (depending on some bits of the current x-coordinate). The current state point is updated unless a distinguished point has been reached (check_dist returns true if (x, y) is a distinguished point). Lines 15 and 16 implement the negation map using a conditional assignment: if the least significant bit of y is 1, the current point is replaced with $-(x, y) = (x, x + y)$.

The Sage code for cycle detection is shown in lines 19 to 25 in Figure 4.1. The variable ctr is a 4-bit counter. To compute the potential exit point of a cyle, we store the x-coordinate of the minimum point (i.e., the point with the smallest

```
def random_step(x, y, ctr, double, x_min):
# point addition/doubling
  T_x = get_precomputed_x(x)
  l1 = x if double else (x + T_x)
  l1_inv = 1/l1
  T_y = get_precomputed_y(x)
  l0 = (x^2 + y) if double else (y + T_y)
  l = l0 * l1_inv
  x3 = l^2 + l + 1
  x3 = x3 if double else x3 + l1
  dist = check_dist(x)
  tmp = l * (x + x3)
  x = x if dist else x3
  y = y if dist else tmp + y + x3
  c_x_y = (get_lsb(y) == 1)
  y = (x + y) if c_x_y else y

# cycle detection
  ctr = (ctr + 1) % 16
  c_ctr = (ctr == 0)
  c_lt = (x < x_min)
  c_new_min = c_lt or c_ctr or double
  double = (x == x_min)
  x_min = x if c_new_min else x_min
  ctr = 0 if c_lt else ctr
  return (x, y, ctr, double, x_min)
```

Fig. 4.1: Sage code for the iteration function.

x-coordinate) of a cycle in x_min. Whenever within 16 steps we reach a smaller point than the current minimum, the flag c_lt is set to true, x_min is updated, and the counter value is reset to 0. A side effect of using the negation map is that we do not need to store the y-coordinate of the minimum point. In case the counter has an overflow to 0, i.e., we did not encounter a cycle within 16 iteration steps, the flag c_ctr is set to true and we move the detection window forward by setting the minimum x_min to the current point. In case we re-visit a point, i.e., the current x-coordinate is equal to x_min, the flag double is set to true for the next iteration resulting in a point doubling in the top part of the code in Figure 4.1. If there just was a point doubling in the current iteration, x_min is updated with the current x-coordinate as well in order to restart cycle detection.

The state of the iteration function consists of the x and y coordinates of the current point. For cycle detection, additionally we require a 4-bit counter, the flag double, and the x-coordinate of the minimum point of the cycle window. For computations in $\mathbb{F}_{2^n}$, in total the state requires $2n+4+1+n$ bits, therefore 344 bits for sect113r2 using $\mathbb{F}_{2^{113}}$ and 386 bits for target117 using $\mathbb{F}_{2^{127}}$.

```
def GF113_inv(x):
    r0 = x^(2^1)
    r1 = r0*x
    r0 = r1^(2^1)
    r1 = r0*x
    r0 = r1^(2^3)
    r1 = r0*r1
    r0 = r1^(2^1)
    r1 = r0*x
    r0 = r1^(2^7)
    r1 = r0*r1
    r0 = r1^(2^14)
    r1 = r0*r1
    r0 = r1^(2^28)
    r1 = r0*r1
    r0 = r1^(2^56)
    r1 = r0*r1
    r0 = r1^(2^1)
    return r0
```

```
def GF127_inv(x):
    r0 = x^(2^1)
    r0 = r0*x
    r0 = r0^(2^1)
    r2 = r0*x
    r0 = r2^(2^3)
    r1 = r0*r2
    r0 = r1^(2^6)
    r1 = r0*r1
    r0 = r1^(2^3)
    r1 = r0*r2
    r0 = r1^(2^15)
    r1 = r0*r1
    r0 = r1^(2^30)
    r1 = r0*r1
    r0 = r1^(2^3)
    r1 = r0*r2
    r0 = r1^(2^63)
    r1 = r0*r1
    r0 = r1^(2^1)
    return r0
```

Fig. 4.2: Sage code for finite-field inversion in $\mathbb{F}_{2^{113}}$ and $\mathbb{F}_{2^{127}}$.

The functional description of the iteration function shows that we need several finite-field operations for the FPGA design, i.e., addition, squaring, multiplication, inversion, and comparison.

4.1. Inversion. Inversion in the finite field is an expensive operation that can be implemented using a sequence of squarings and multiplications to compute $a^{-1} = a^{2^n-2}$. Figure 4.2 shows the inversion ladders that we are using. The shortest addition chain for 112 requires 8 additions, the chain for 126 requires 9 additions [16]. This allows us to compute short addition chains for $2^{113}-2$ and $2^{127}-2$. The inversion procedures require 8 multiplications and 112 squarings for $\mathbb{F}_{2^{113}}$, 9 multiplications and 126 squarings for $\mathbb{F}_{2^{127}}$.

Consecutive squarings can be combined to powers of higher order depending on which power is suitable and most efficient for the implementation.

4.2. Low level functions. The main components for implementing logical expressions on an FPGA are lookup tables (LUTs). The LUTs in the Spartan-6 are LUT-6 with 6 input wires. However, internally each LUT-6 is implemented with two LUT-5 using the same input wires. The sixth input wire selects the final output by controlling a 2-bit multiplexer (see Figure 4.3). The LUTs can be configured either as LUT_6 providing LUT-6 functionality or as LUT_6_2 that gives access to each output of the two LUT-5.

4.3. Addition. This operation requires a very small amount of logic. $\lceil k/2 \rceil$ LUTs in LUT_6_2 configuration (providing $\geq k$ output bits) are sufficient

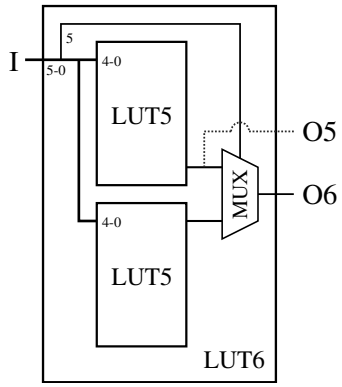

Fig. 4.3: Structure of a LUT-6 implemented as two LUT-5. Input I is a 6-bit bus. Output O5 is available when used as LUT_6_2.

for the implementation. However, often addition can be combined with follow-up operations like squaring such that the logic might be absorbed. We do not explicitly implement addition with LUT_6_2 components but leave it to the Xilinx tool chain to map the VHDL code.

4.4. Squaring. This operation requires simply inserting zeros between the coefficients and performing a reduction modulo the irreducible polynomial chosen for the respective field (which is a trinomial for both our targets). Consecutive squarings can be combined in order to absorb logic into a smaller number of LUTs. Single squarings appear in combination with addition in our design. Therefore, we express single squarings as VHDL code and leave it to the Xilinx tool chain to combine the logic. For sequences from 2 to 8 squarings, we generate optimized logic and explicitly use LUT_6_2 components.

4.5. Multiplication. This operation is the most expensive operation in terms of area. We use three levels of Karatsuba multiplication for both fields. Since we are operating on operands with a prime number of bits, we cannot apply Karatsuba straight away. For operations in $\mathbb{F}_{2^{113}}$, we handle the top bits of the operands separately and perform Karatsuba on 112-bit operands, resulting in 27 multiplications of 14-bit polynomials. For operations in $\mathbb{F}_{2^{127}}$, we simply add a zero bit as most significant bit to each operand and perform 27 multiplications of 16-bit polynomials at the cost of a small overhead.

For the low-level multiplications, we generate optimized logic using LUT_6_2 components. Figure 4.4 shows an example of how we cover the terms; adding up the columns in the figure gives the result of a 7×7 polynomial multiplication. Using a LUT-6 as in the dashed box covers only three terms using 6 inputs; using a LUT_6_2 as in the solid box requires only five separate inputs but covers 4 terms, the two output wires of the LUT_6_2 are used for the two involved columns (requiring independent sums). Special care needs to be taken at the boundaries. Additional logic is required to sum up over each column.

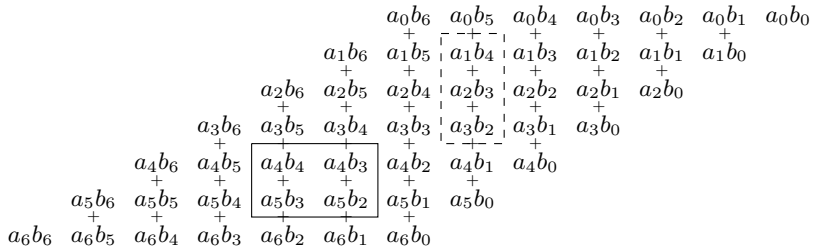

Fig. 4.4: Example for assigning terms of a 7×7 polynomial multiplication to LUT-6's. The dashed box requires six inputs but covers only three terms. The solid box requires only five inputs for four terms and thus can be implemented using a LUT_6_2.

Also for the preparation of the inputs for the low-level multiplications and for the computation of the total result we are generating optimized logic using LUT_6_2 components whenever possible. All in all, one $\mathbb{F}_{2^{113}} \times \mathbb{F}_{2^{113}}$ multiplication requires on average 3071 LUTs and one $\mathbb{F}_{2^{127}} \times \mathbb{F}_{2^{127}}$ multiplication on average 3620 LUTs (after placement and routing). The implementation of the multiplication is pipelined and requires three clock cycles.

4.6. Comparison. For cycle detection we require a less-than comparison and an equality check on the same inputs. We implemented optimized logic to compute both operations at once using LUT_6_2 components.

4.7. Implementing the iteration function. Our goals are high throughput with low overhead. Therefore, we implement the main part of the iteration function with fully pipelined, unrolled code. All components are busy all the time: the design computes one step of the random walk in each cycle while working on many independent random walks in parallel in a pipelined fashion.

However, using this approach for finite-field inversion as well would require a large amount of resources. The iteration function requires two multiplications, one inversion, and various additions and squarings. The inversion itself requires 8 multiplications, thus demanding more than 80% of the total resources. We do better by using Montgomery's trick for inversion, combining n inversions into $3(n-1)$ multiplications and just 1 inversion.

We use a dual-buffer design to implement a pipelined version of the Montgomery inversion. The buffers are used as follows:

1. Fill buffer 1 with data, using one multiplier to compute the overall product of the buffer.
2. Invert the product of buffer 1 while filling buffer 2. The size of the buffers must be large enough to hide the latency of one inversion.
3. Once the inverse of the first product has been computed, empty buffer 1 in reverse order by computing the individual inverses using two multipliers

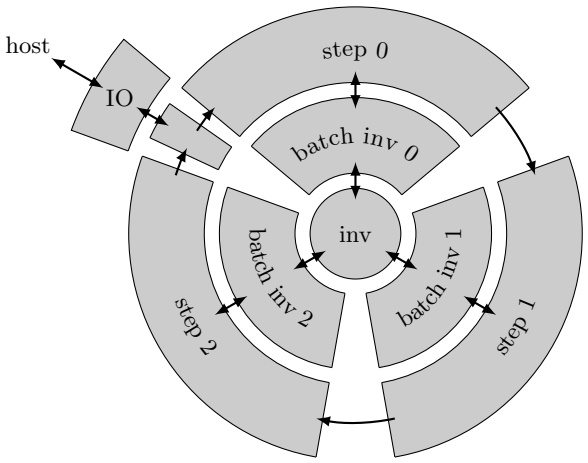

Fig. 4.5: Overall layout of the design with one single inverter (inv) and three unrolled iteration functions (step 1-3) using Montgomery's trick for inversion (batch inv 1-3).

while at the same time filling buffer 1 again using the just emptied slots. At the same time, invert the overall product of buffer 2.

4. Continue iteratively by filling buffer 1 and buffer 2 periodically, alternating in ascending and descending order.

Given a sufficient amount of buffer memory, the latency of the actual inversion within the Montgomery inversion does not matter. Therefore, we implemented the inversion not as unrolled code, but as an application specific instruction set processor (ASIP) with a custom-made instruction set and one single finite-field multiplier. This results in a lower area consumption than when using an unrolled core for the price of a higher overall latency and lower utilization of the logic for the finite-field operations within the inverter.

We do even better by sharing one inverter between several instances of the iteration function, giving space for more instances. The additional operations (multiplications in the aggregation and crossmultiplications; logic) for Montgomery's trick are simply implemented on the inversion core without cost for further multipliers. This increases the overall latency of the inversion, but as explained above we use buffering to hide this latency.

To simplify development and to increase flexibility, we wrote tools that automatically generate unrolled code (for the iteration function) and an ASIP (for the inversion) from Sage code. Therefore, the logic of the iteration function can easily be tested using Sage and also easily be altered at any time. For each of our target curves, we are able to put three instances of the iteration function on one Spartan-6 XC6SLX150 FPGA, whereas using separate inverters for each core (as in [15]) would allow at best two instances.

Module	Inst.	Mult. per instance	$\mathbb{F}_{2^{113}}$ LUTs per instance	$\mathbb{F}_{2^{113}}$ LUTs total	$\mathbb{F}_{2^{113}}$ FPGA util.	$\mathbb{F}_{2^{127}}$ LUTs per instance	$\mathbb{F}_{2^{127}}$ LUTs total	$\mathbb{F}_{2^{127}}$ FPGA util.
3 cores with inv.				57,956	63%		68,376	74%
iteration func.	3	2	7,789	23,368	25%	9,123	27,369	30%
$\mathbb{F}_{2^{113}}$ inv.	1	1	5,817	5,817	6%	7,151	7,151	8%
batching	3	3	9,535	28,605	31%	11,397	34,191	37%
$\mathbb{F}_{2^{113}}$ mult.	16		3,071	49,131	53%	3,620	57,927	64%
total (incl. IO)				63,388	69%		72,919	79%

Fig. 4.6: Area consumption by component. All values are post placement.

4.8. Overall architecture. To streamline the design, we arrange the instances of the iteration function in a circle: the output of one instance is the input for the next one. Therefore, the design requires only one single IO point. Figure 4.5 shows the overall layout of our design.

The host computer randomly computes starting points of independent random walks using a 64-bit seed. During computation of the random walks on the FPGAs, the seeds are stored on the host computer; only a 12-bit temporal ID is sent to the FPGAs along with the point coordinates in order to associate random walks with their seeds.

Each IO instance stores incoming data in a buffer. If there is an empty slot in the pipeline (either during setup phase in the beginning or because a distinguished point has been returned to the host computer), the IO interface sets up the state of a new random walk using a fresh starting point and its 12-bit ID and puts it into the pipeline. Every clock cycle, the pipeline feeds a state into the first instance of the iteration function (step 0) which computes one step of the random walk. During the inversion, the state data of the random walk is stored in a buffer (batch inv 0). Once the first iteration step is computed, the pipeline forwards the state to the next instance of the iteration function (step 1). The state circles through the instances, step by step computing a random walk, until a distinguished point is reached.

Now, the distinguished point is forwarded back to the IO interface. Random-walk computations are still performed on the state on the way through the instances, but the state is not updated anymore (see lines 13 and 14 in Figure 4.1). Once a state with a distinguished point arrives at the IO interface, the interface returns the x-coordinate and the ID of the distinguished point to the host and fills the pipeline slot using a fresh input point. The host associates the original 64-bit seed with the distinguished point using the 12-bit ID and sends the seed and the x-coordinate to the server which sorts incoming points, detects collisions, and finally computes the discrete logarithm.

Table 4.6 shows the area demand of the design for the different components for both fields.

The final 3-core design routes for and runs at 100MHz for both finite-field implementations. For the $\mathbb{F}_{2^{113}}$ case, we also tested a design with 4 iteration

cores on one XC6SLX150 FPGA. However, the power consumption of the design was too high and the design did not run stably, producing incorrect results. Furthermore, we tried to increase the frequency of the design by introducing additional pipelining steps and using different routing strategies. We were able to place and route and also run a single-core design at up to 160MHz; however, our attempts to increase the frequency of the 3-core design failed to produce stable, operational designs, because the power consumption at these frequencies was too high.

For testing our designs we used Spartan-6 development boards from SciEngines and from Opal Kelly. For running the attacks, we used two "Rivyera" FPGA-cluster computers from SciEngines with 64 Spartan-6 FPGAs each. A Rivyera is a classical "host" computer combined with to up to 128 FPGAs (in the high-density version up to 256 FPGAs). The FPGAs are connected to the host computer using a PCIe host interface. SciEngines provides an API for programming the FPGAs and for communication between the FPGAs and between FPGAs and the host computer.

4.9. Solving DLPs. We solved a 112-bit DLP on the sect113r2 curve with an earlier 2-core, 100MHz design within about 48.1 days using up to 120 Spartan-6 FPGAs. (Not all FPGAs were available for our computations all the time.) The solution involved the computation of 82,177,699 distinguished points. The expected duration was 30.8 days for the computation of $\sqrt{\pi \cdot 2^{112}/4}/2^{30} \approx$ 59,473,682 distinguished points.

We are currently doing computations for solving a 117.35-bit DLP on the elliptic curve target117 over $\mathbb{F}_{2^{127}}$. We are using the same property for distinguished points as for our sect113r2 computations, namely the top 30 bits from the x-coordinate being zero. Therefore, we expect to require $\sqrt{\pi \cdot 2^{117.35}/4}/2^{30} \approx$ 379,821,956 distinguished points for solving the DLP. Using all 128 FPGAs of our cluster, we expect the computation to be finished after about 123 days.

4.10. Power consumption. We measured the power consumption of our designs in the following way: One of our Rivyera FPGA clusters with 64 FPGAs requires about 215W when the FPGAs are not programmed. Under full usage, while our $\mathbb{F}_{2^{113}}$ design is running, the total power demand is about 725W. Thus, 64 FPGAs require about 510W while running the $\mathbb{F}_{2^{113}}$ design; a single FPGA requires about 8W. For the $\mathbb{F}_{2^{127}}$ design the total power consumption is 755W. Therefore, 64 FPGAs require about 540W, a single FPGA about 8.4W.

5 Experiments

The obvious way to verify the performance and functionality of our implementation is to repeat the following procedure many times: generate a random point Q on the curve sect113r2, use the implementation to find k such that $Q = kP$, see how long this takes, and check that in fact $Q = kP$.

The reason for repeating this procedure many times is that the performance is a random variable. Checking the performance of a *single* DL computation

would obviously be inadequate as a verification tool. For example, if the claimed *average* DL time is T while the observed time of a single DL computation is $2.3T$, then it could be that this particular computation was moderately unlucky, or it could be that the claim was highly inaccurate.

There are two reasons that more efficient verification procedures are important. First, it was feasible for us to carry out a sect113r2 DL computation, but performing *many* such computations would have been quite expensive. Second, and more importantly, verification is not merely something to carry out in retrospect: it provides essential feedback during the exploration of the design space. Below we describe the verification steps that we took for our final implementation, but there were also many rounds of similar verification steps for earlier versions of the implementation.

Running hundreds or thousands of walks (a tiny fraction of a complete sect113r2 DL computation; recall that we expect orders of magnitude more distinguished points for our selected parameters) produces reasonably robust statistics regarding the number of iterations required to find a distinguished point, and regarding the time used for each iteration. However, it does not provide any evidence regarding the number of distinguished points required to compute a DL. A recurring theme of several recent papers is that standard heuristics overestimate the randomness of DL walks, and thus underestimate the number of distinguished points required; see, e.g., the correction factors in [1, Appendix B] and the further correction factors in [5, Section 4].

To efficiently verify performance *including* walk randomness and successful DL computation, we adapt the following observation from Bernstein, Lange, and Schwabe [6]. The fastest available ECDL algorithms use the fastest available formulas for adding affine points, and those formulas are independent of some of the curve coefficients: specifically, [6] used formulas that are independent of b in $y^2 = x^3 - 3x + b$, and we use formulas that are independent of b in $y^2 + xy = x^3 + ax^2 + b$. The same algorithms thus work without change for points (and precomputed tables) on other curves obtained by varying b. Searching many curves finds curves with different sizes of prime-order subgroups, allowing tests of exactly the same ECDL algorithms at different scales.

For example, applying an isomorphism to sect113r2 to obtain $a = 1$ as described earlier, and then changing b to 10010111, produces a curve with a subgroup of prime order $1862589870449786557 \approx 2^{60.69}$. This group is large enough to carry out reasonably large experiments without distractions such as frequent self-colliding walks, and at the same time small enough for experiments to complete quickly.

We performed 1024 DL computations on this curve, in each case using 20 bits to define distinguished points. These computations used a total of 1201100 walks. The average number of walks per DL was slightly over 1173. For comparison, the predicted average is $\sqrt{\pi\ell/4}/2^{20} \approx 1153$ for $\ell = 1862589870449786557$, and the predicted standard deviation is on the same scale as the predicted average; the gap between 1173 and 1153 is unsurprising for 1024 experiments. Each computation successfully produced a verified discrete logarithm.

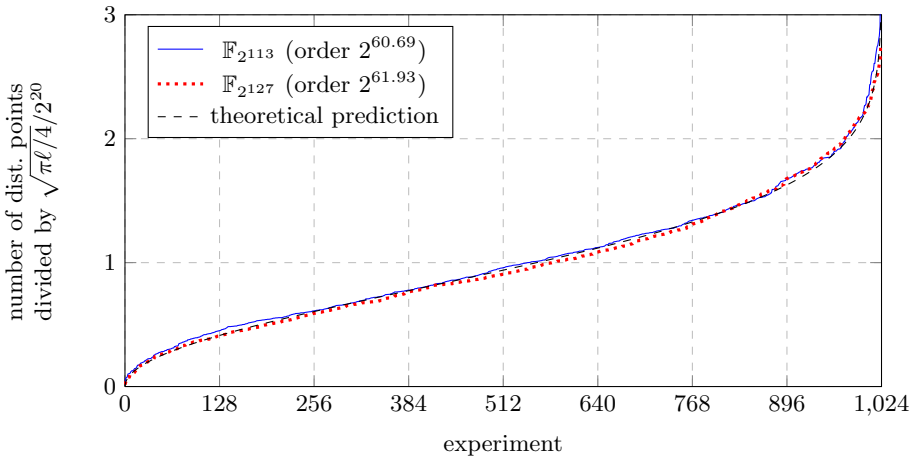

Fig. 5.1: Blue curve: 1024 independent experiments for $\mathbb{F}_{2^{113}}$. The experiments are sorted by the number of distinguished points they required to find a discrete logarithm, and are then placed at $x = 0$, $x = 1$, ..., $x = 1023$ respectively. The y-axis is the number of distinguished points divided by $\sqrt{\pi\ell/4}/2^{20}$. Dotted red curve: 1024 independent experiments for $\mathbb{F}_{2^{127}}$, organized the same way. Dashed black curve: $y = \sqrt{-(4/\pi)\log(1 - x/1024)}$ from standard rho theory. This curve shows that our experiments are close to the expected distribution.

We defined the first DL computation to use seed $0, 1, 2, \ldots$ until finding a collision between seed s and a previous seed; the second DL computation to use seeds $s + 1, s + 2, \ldots$ until finding a collision within those seeds; etc. We post-processed seeds with AES before multiplying them by Q, so (if AES is strong) choosing consecutive seeds is indistinguishable from choosing independent uniform random 128-bit scalars.

The advantage of choosing consecutive seeds is that, without knowing in advance which seeds would be used in each computation, we simply provided a large enough batch of seeds $0, 1, 2, \ldots$ to our FPGAs. Retroactively attaching each seed to the correct computation was a simple matter of sorting the resulting distinguished points in order of seeds and then scanning for collisions. The sorting step here is important: if we had scanned for collisions using the order of points output by the FPGAs then we would have incorrectly biased the initial computations towards short walks.

We performed the same experiment on a curve over $\mathbb{F}_{2^{127}}$ with a subgroup of order $\ell_2 = 43898321882824425 01 \approx 2^{61.93}$. In this case 1024 DL computations required 1,792,905 distinguished points. The average number of distinguished points per DL was 1751. The predicted average is $\sqrt{\pi\ell_2/4}/2^{20} \approx 1771$. Figure 5.1 shows the number of required distinguished points divided by the predicted average for both experiments. The experiments have been ordered by the number of required distinguished points.

We also carried out various experiments over $\mathbb{F}_{2^{113}}$ with

- a group of size $214943357179500410153 9 \approx 2^{70.86}$ with $b = \texttt{110}$,
- a group of size $2608103394926752635062767 \approx 2^{81.11}$ with $b = \texttt{100111}$, and
- a group of size $153412233055515912111528877 7 \approx 2^{90.31}$ with $b = \texttt{10000111}$.

We spot-checked walks against a separate software implementation, verified correctness of 16 DL computations for the 70-bit group, and verified correctness of 1 DL computation for the 80-bit group.

6 Comparison

Section 1 summarized how our results improve upon the recent results [32] from Wenger and Wolfger. This section compares the results and techniques of our 113-bit implementation with theirs in more detail.

6.1. Multiplier area. A few different multiplier structures are considered in [32, Section 6.8 and Appendix B]. The best results, 3757 LUTs, rely on traditional power-of-2 Karatsuba multipliers, for example building a 64-bit multiplier from three 32-bit multipliers. We do much better, 3071 LUTs, by exploring a much wider range of optimizations: in particular, we drop the power-of-2 restriction, allowing efficient use of three Karatsuba levels, and we optimize the low-level usage of LUTs. These optimizations should also be useful for constructive applications.

6.2. Number of multipliers. The design in [32] applies an inverter to a batch of inputs, where each core provides one input from its first walk; then applies the same inverter to another batch of inputs, where each core provides one input from its second walk; etc. This means that the batch size is very small even on the bigger FPGA and would become ridiculous (size two) if mapped to our FPGA. This requires a high-throughput inverter: [32] uses a fully unrolled inverter, requiring 8 multipliers and 112 squarers.

We instead use a dual-buffer memory to batch inversions across cores *and* across many random walks from each core. This lets us use a high-latency inverter without slowing down the rest of the design. This, in turn, allows us to use a low-area ASIP design for the inverter, requiring only one multiplier and one module each to compute a^2, a^{2^2}, a^{2^4}, and a^{2^8}.

A slight disadvantage is that for c cores we need $3c$ multipliers to batch inversions, whereas in [32] one core can skip 3 of these multipliers, for a total of just $3c - 3$ multipliers to batch inversions. Furthermore, to simplify routing we synthesized 5-core, 6-core, 7-core, and 8-core designs as two separate clusters, with a separate inverter in each cluster. However, overall we still save the area for 3 multipliers in $\leq$8-core designs, or 4 multipliers in $\leq$4-core designs.

6.3. Total area. All in all, our improvements and optimizations reduce the area cost significantly compared to [32]. As noted in Section 1, we do not have access to any Kintex-7 FPGAs for testing, but for comparability we nevertheless

cores	LUTs	registers	RAMs	slices	source
5	73%	41%	31%	80%	[32]
6	62%	28%	15%	84%	new
7	71%	33%	17%	91%	new
8	80%	37%	19%	96%	new

Fig. 6.1: Resource utilization of several ECDL designs synthesized for the XC7K325T-2 at 180MHz.

synthesized our design for the XC7K325T-2 used in the KC705 development boards in [32]. Table 6.1 shows that our 7-core design uses fewer LUTs, fewer registers, and fewer block RAMs than the 5-core design in [32]; it uses more slices but we expect it to run stably at the same 180MHz. We assume that [32] synthesized the design for larger frequencies (resulting in a closer placement and thus requiring fewer slices) and later experimentally tried out the maximum frequency delivering stable results. Note that overheating is the main bottleneck identified in [32], and heat is generated primarily by computation and memory access, not by chip area per se. With access to this FPGA we could verify stability for 7 or 8 cores and possibly fine-tune the design to allow higher frequencies.

6.4. Fruitless cycles. A further advantage of our design is that we waste fewer iterations on fruitless cycles. Specifically, we use doublings to escape fruitless cycles, while [32] uses additions. The detailed analysis in [9] indicates that additions create new types of fruitless cycles, whereas doublings avoid this problem.

[32, p. 4] argues that using additions to escape fruitless cycles has "a huge advantage when a hardware design is done" since "no on-chip point doubling circuit is necessary". However, in our unrolled design, the doubling circuit reuses the addition circuit with miniscule additional area cost.

We performed many experiments to check that our iteration function is correctly computed by our implementation and that it is as effective as expected, gaining a factor $\sqrt{2}$ in the average observed number of iterations compared to not using the negation map; see Section 5. The small-scale negation-map experiments reported in [32, Table 2] show a speedup factor only 1.32, i.e., 6% worse than $\sqrt{2}$. This gap is consistent with the analysis in [9].

[32, Table 2, "Point doubling" entry] reports doubling experiments that were also 6% worse than $\sqrt{2}$. This is inconsistent with the analysis and experiments in [9] (and with our own experiments); [32] does not discuss this inconsistency. Presumably the "Point doubling" experiment in [32] actually used a slightly different cycle-escape method from what we call doubling, but no further details are provided in [32].

6.5. Target curves. [32] illustrates its techniques by attacking the SECG curve sect113r1, while we illustrate our techniques by attacking the SECG curve

sect113r2 and our new curve target117. The prime orders are

$$519229685853482768983588257883070 3 \approx 2^{112.00000000000000001703}$$
$$519229685853482770297249790995240 3 \approx 2^{112.00000000000000002068}$$
$$212146114040485326348618959071598183 \approx 2^{117.35254157354507970215}$$

respectively. The sect113r1 group size was summarized as "2^{112}" in [32] but as "bit size 113" in [17]. Note that rounding the exponent to the nearest integer also produces 112 for the earlier ECDL record in [7], whereas rounding the exponent up creates a large separation (41% difference in estimated attack cost) between practically identical orders marginally above and below a power of 2. We suggest instead rounding to two digits after the decimal point in the exponent: $2^{111.78}$ for [7], $2^{112.00}$ for [32], $2^{112.00}$ for sect113r2 in this paper, and $2^{117.35}$ for target117 in this paper.

6.6. Generality and novelty of results. In summary, [32] implements only one curve and extrapolates to larger sizes, while we implemented two different field sizes and performed many more DL computations.

We have designed a considerably smaller multiplier and an iteration function that consumes significantly fewer LUTs. The design from [32] cannot run on our FPGAs and quick adjustments are not possible because of their huge inverter.

Reasons for the size difference are our improvements in

- Designing a smaller multiplier, using more levels of Karatsuba and no power-of-2 restriction.
- Designing a low area inverter with bigger batch size and a different way to batch across iterations.
- A new definition of $|P|$, saving $n-1$ bit comparisons.
- Rather than computing coefficients in every walk, compute coefficients only for the two colliding walks (negligible cost at the end of the DL). This saves all circuitry for computations $\mod l$.
- A new way of checking for cycles, reducing cost to one field element and a 4-bit counter.
- Integrating the doubling circuitry with the general addition circuitry, removing the overhead for dealing with fruitless cycles.

Furthermore we use doublings to escape fruitless cycles, so that we decrease the overhead and avoid using the bottom bit in deciding the next step.

References

1. D. V. Bailey, L. Batina, D. J. Bernstein, P. Birkner, J. W. Bos, H.-C. Chen, C.-M. Cheng, G. V. Damme, G. de Meulenaer, L. J. D. Perez, J. Fan, T. Güneysu, F. Gürkaynak, T. Kleinjung, T. Lange, N. Mentens, R. Niederhagen, C. Paar, F. Regazzoni, P. Schwabe, L. Uhsadel, A. V. Herrewege, and B.-Y. Yang. Breaking ECC2K-130. Cryptology ePrint Archive, Report 2009/514, 2009. https://eprint.iacr.org/2009/541/. 2, 10, 20

2. D. J. Bernstein. Newton's identities without factorization, 1997. https://cr.yp.to/papers/newton.pdf. 9
3. D. J. Bernstein, H. Chen, C. Cheng, T. Lange, R. Niederhagen, P. Schwabe, and B. Yang. ECC2K-130 on NVIDIA GPUs. In G. Gong and K. C. Gupta, editors, *Progress in Cryptology - INDOCRYPT 2010*, volume 6498 of *Lecture Notes in Computer Science*, pages 328–346. Springer, 2010. https://cr.yp.to/papers.html#ecc2k130. 2
4. D. J. Bernstein and T. Lange. Computing small discrete logarithms faster. In S. D. Galbraith and M. Nandi, editors, *Progress in Cryptology - INDOCRYPT 2012: 13th International Conference in Cryptology in India*, volume 7668 of *Lecture Notes in Computer Science*, pages 317–338, Kolkata, India, Dec. 9–12, 2012. Springer, Heidelberg, Germany. https://eprint.iacr.org/2012/458. 5
5. D. J. Bernstein and T. Lange. Two grumpy giants and a baby. In E. W. Howe and K. S. Kedlaya, editors, *ANTS X: proceedings of the tenth algorithmic number theory symposium*, pages 87–111. Mathematical Sciences Publishers, 2013. https://eprint.iacr.org/2012/294. 20
6. D. J. Bernstein, T. Lange, and P. Schwabe. On the correct use of the negation map in the Pollard rho method. In D. Catalano, N. Fazio, R. Gennaro, and A. Nicolosi, editors, *Public Key Cryptography – PKC 2011*, volume 6571 of *LNCS*, pages 128–146. Springer, 2011. https://cryptojedi.org/papers/#negation. 10, 11, 20
7. J. W. Bos, M. E. Kaihara, T. Kleinjung, A. K. Lenstra, and P. L. Montgomery. Playstation 3 computing breaks 2^{60} barrier; 112-bit prime ECDLP solved, 2009. http://lacal.epfl.ch/112bit_prime. 2, 10, 24
8. J. W. Bos, M. E. Kaihara, and P. L. Montgomery. Pollard rho on the PlayStation 3. In *Workshop Record of SHARCS'09: Special-purpose Hardware for Attacking Cryptographic Systems*, pages 35–50, 2009. https://hyperelliptic.org/tanja/SHARCS/record2.pdf. 2, 10
9. J. W. Bos, T. Kleinjung, and A. K. Lenstra. On the use of the negation map in the Pollard rho method. In G. Hanrot, F. Morain, and E. Thomé, editors, *Algebraic Number Theory*, volume 6197 of *LNCS*, pages 66–82. Springer, 2010. http://www.joppebos.com/files/negation.pdf. 11, 23
10. Certicom Research. SEC 2: Recommended elliptic curve domain parameters, version 1.0, 2000. http://www.secg.org/SEC2-Ver-1.0.pdf. 4
11. Certicom Research. SEC 2: Recommended elliptic curve domain parameters, version 2.0, 2010. http://www.secg.org/sec2-v2.pdf. 4
12. I. M. Duursma, P. Gaudry, and F. Morain. Speeding up the discrete log computation on curves with automorphisms. In K. Lam, E. Okamoto, and C. Xing, editors, *Advances in Cryptology - ASIACRYPT '99, International Conference on the Theory and Applications of Cryptology and Information Security, Singapore, November 14-18, 1999, Proceedings*, volume 1716 of *Lecture Notes in Computer Science*, pages 103–121. Springer, 1999. 27, 28, 29
13. ECC Brainpool. ECC Brainpool standard curves and curve generation, 2005. http://www.ecc-brainpool.org/download/Domain-parameters.pdf. 6
14. S. Engels. Breaking ecc2-113: Efficient implementation of an optimized attack on a reconfigurable hardware cluster, 2014. http://www.emsec.rub.de/media/attachments/files/2014/11/MA_Engels.pdf. 2
15. J. Fan, D. V. Bailey, L. Batina, T. Güneysu, C. Paar, and I. Verbauwhede. Breaking elliptic curve cryptosystems using reconfigurable hardware. In *International Conference on Field Programmable Logic and Applications, FPL 2010*, pages 133–138. IEEE, 2010. https://www.cs.ru.nl/~lejla/FPL2010.pdf. 2, 17

16. A. Flammenkamp. Shortest addition chains, 2008. http://wwwhomes.uni-bielefeld.de/achim/addition_chain.html. 14
17. S. D. Galbraith and P. Gaudry. Recent progress on the elliptic curve discrete logarithm problem. *Des. Codes Cryptography*, 78(1):51–72, 2016. https://eprint.iacr.org/2015/1022. 6, 24
18. R. P. Gallant, R. J. Lambert, and S. A. Vanstone. Improving the parallelized Pollard lambda search on anomalous binary curves. *Mathematics of Computation*, 69(232):1699–1705, 2000. 6
19. R. J. Harley. Solution to Certicom's ECC2K-95 problem (email message), 1998. http://cristal.inria.fr/~harley/ecdl5/ECC2K-95.submission.text. 10
20. Y. Hitchcock, P. Montague, G. Carter, and E. Dawson. The security of fixed versus random elliptic curves in cryptography. In R. Safavi-Naini and J. Seberry, editors, *ACISP 03: 8th Australasian Conference on Information Security and Privacy*, volume 2727 of *Lecture Notes in Computer Science*, pages 55–66, Wollongong, NSW, Australia, July 9–11, 2003. Springer, Heidelberg, Germany. 5
21. F. Kuhn and R. Struik. Random walks revisited: Extensions of Pollard's rho algorithm for computing multiple discrete logarithms. In S. Vaudenay and A. M. Youssef, editors, *SAC 2001: 8th Annual International Workshop on Selected Areas in Cryptography*, volume 2259 of *Lecture Notes in Computer Science*, pages 212–229, Toronto, Ontario, Canada, Aug. 16–17, 2001. Springer, Heidelberg, Germany. http://grouper.ieee.org/groups/802/PrivRecsg/email/pdfJ3n8ucpxL8.pdf. 5
22. National Institute for Standards and Technology (NIST). Recommended elliptic curves for Federal government use, 1999. http://csrc.nist.gov/groups/ST/toolkit/documents/dss/NISTReCur.pdf. 5
23. National Institute for Standards and Technology (NIST). Digital signature standard (DSS) FIPS 186–4, 2013. http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.186-4.pdf. 5
24. National Security Agency. Suite B cryptography / cryptographic interoperability, 2005. http://www.nsa.gov/ia/programs/suiteb_cryptography/. 6
25. T. Oliveira, J. López, D. F. Aranha, and F. Rodríguez-Henríquez. Lambda coordinates for binary elliptic curves. In G. Bertoni and J.-S. Coron, editors, *CHES*, volume 8086 of *LNCS*, pages 311–330. Springer, 2013. https://eprint.iacr.org/2013/131. 6
26. J. M. Pollard. Monte Carlo methods for index computation (mod p). *Mathematics of Computation*, 32(143):918–924, 1978. http://www.ams.org/journals/mcom/1978-32-143/S0025-5718-1978-0491431-9/S0025-5718-1978-0491431-9.pdf. 9
27. G. Seroussi. Compact representation of elliptic curve points over $\mathbb{F}_{2^n}$. *HP Labs Technical Reports*, HPL-98-94R1, 1998. http://www.hpl.hp.com/techreports/98/HPL-98-94R1.html. 9
28. N. J. A. Sloane. The on-line encyclopedia of integer sequences, 2016. https://oeis.org. 29
29. E. Teske. On random walks for Pollard's rho method. *Mathematics of Computation*, 70(234):809–825, 2001. http://www.ams.org/journals/mcom/2001-70-234/S0025-5718-00-01213-8/S0025-5718-00-01213-8.pdf. 10
30. P. C. van Oorschot and M. J. Wiener. Parallel collision search with cryptanalytic applications. *Journal of Cryptology*, 12(1):1–28, 1999. http://www.scs.carleton.ca/~paulv/papers/JoC97.pdf. 9
31. E. Wenger and P. Wolfger. Solving the discrete logarithm of a 113-bit Koblitz curve with an FPGA cluster. In A. Joux and A. Youssef, editors, *Selected Areas in Cryptography – SAC 2014*, volume 8781 of *LNCS*, pages 363–379. Springer, 2014. https://eprint.iacr.org/2014/368. 2, 10

32. E. Wenger and P. Wolfger. Harder, better, faster, stronger: elliptic curve discrete logarithm computations on FPGAs. *Journal of Cryptographic Engineering*, pages 1–11, 2015. https://eprint.iacr.org/2015/143. 2, 3, 4, 5, 10, 11, 22, 23, 24
33. M. J. Wiener and R. J. Zuccherato. Faster attacks on elliptic curve cryptosystems. In S. E. Tavares and H. Meijer, editors, *Selected Areas in Cryptography '98, SAC'98, Kingston, Ontario, Canada, August 17-18, 1998, Proceedings*, volume 1556 of *Lecture Notes in Computer Science*, pages 190–200. Springer, 1998. 6

A Parameter for transforming the curve

$$t = w^{112} + w^{111} + w^{108} + w^{107} + w^{106} + w^{104} + w^{101} + w^{96} + w^{95} + w^{91} + w^{89} + w^{88} + w^{87} + w^{86} + w^{83} + w^{82} + w^{81} + w^{80} + w^{78} + w^{75} + w^{74} + w^{67} + w^{64} + w^{63} + w^{62} + w^{61} + w^{60} + w^{58} + w^{57} + w^{53} + w^{50} + w^{49} + w^{46} + w^{43} + w^{42} + w^{41} + w^{39} + w^{37} + w^{36} + w^{33} + w^{32} + w^{31} + w^{30} + w^{28} + w^{26} + w^{24} + w^{23} + w^{19} + w^{17} + w^{15} + w^{14} + w^{13} + w^{11} + w^{10} + w^9 + w^7 + w^5 + w^4 + w^3 + w^2$$

Hexadecimal representation of t: `0x19D218BCF4C09F6264EB3D58AEEBC`.

B Probabilities of fruitless cycles

The following model of fruitless cycles is implicit in [12, Proposition 3.1]. Let $T_0, \ldots, T_{n-1}$ be basis vectors of an n-dimensional lattice. For each lattice point R choose an independent uniform random $h(R) \in \{T_0, \ldots, T_{n-1}, -T_0, \ldots, -T_{n-1}\}$. Starting from a lattice point R_0, define $R_1 = R_0 + h(R_0)$, $R_2 = R_1 + h(R_1)$, etc.

Evidently $R_2 = R_0$ if and only if $h(R_1) = -h(R_0)$. This fruitless 2-cycle occurs with probability exactly $\delta/2$ where $\delta = 1/n$.

Even if all 2-cycles are caught and eliminated, larger cycles can occur. For example, for any $n \geq 2$, one can have a fruitless 4-cycle $R_4 = R_0$ with distinct R_0, R_1, R_2, R_3. This occurs if, e.g., $h(R_0) = T_0$, $h(R_1) = T_1$, $h(R_2) = -T_0$, and $h(R_3) = -T_1$.

It is often stated that larger cycle lengths occur less frequently. Specifically, [12] shows for each $t \in \{2, 4, 6, 8, \ldots\}$ that a cycle of length exactly t starting from R_0 occurs with probability $O(\delta^{t/2})$. However, the constant implicit in O increases quite rapidly with t. For example, taking $n = 1024$ (as in our computation) and $t = 16$ means that $\delta = 2^{-10}$ and $\delta^{t/2} = 2^{-80}$, but the bound in [12] is much larger, approximately 2^{-40}. If this bound is tight then fruitless cycles are a problem for us: even if we detect and eliminate all cycle lengths smaller than 16, we would expect 16-cycles to appear approximately once every 2^{40} steps.

Fortunately, the bound is far from tight. We have, with computer assistance, computed the exact probabilities of fruitless t-cycles in this model for several small values of t, and checked several of these formulas against the results of a

comprehensive simulation for $n = 10$:

t	probability
2	$\frac{1}{2}\delta$
4	$\frac{1}{4}\delta^2 - \frac{1}{4}\delta^3$
6	$\frac{1}{2}\delta^3 - \frac{21}{16}\delta^4 + \frac{13}{16}\delta^5$
8	$\frac{27}{16}\delta^4 - \frac{131}{16}\delta^5 + \frac{415}{32}\delta^6 - \frac{207}{32}\delta^7$
10	$\frac{31}{4}\delta^5 - \frac{3755}{64}\delta^6 + \frac{10615}{64}\delta^7 - \frac{25965}{128}\delta^8 + \frac{11253}{128}\delta^9$
12	$\frac{1415}{32}\delta^6 - \frac{60795}{128}\delta^7 + \frac{524985}{256}\delta^8 - \frac{2242229}{512}\delta^9 + \frac{2322943}{512}\delta^{10} - \frac{14221}{8}\delta^{11}$
14	$\frac{4779}{16}\delta^7 - \frac{274463}{64}\delta^8 + \frac{6638303}{256}\delta^9 - \frac{42763903}{512}\delta^{10} + \frac{305264211}{2048}\delta^{11} - \frac{35158655}{256}\delta^{12} + \frac{102125321}{2048}\delta^{13}$

Fix t. We compute the probability for t as follows, generalizing the obvious $t = 2$ computation, and generalizing the $t = 4$ computation that appeared in [12, Proposition 3.1, case "$j_1 = j_3, j_2 = j_4$"].

Consider t-step self-avoiding closed paths in $\mathbb{Z}^k$ starting at $(0, 0, \ldots, 0)$, where each step is adding or subtracting 1 in a single coordinate. "Closed" means that the path ends at its starting point. "Self-avoiding" means that the only collision in the path is the collision of the starting point and the ending point.

Define p_k as the number of such paths that use all k coordinates. Note that each of the k coordinates must be used in at least two steps, so $p_k = 0$ if $k > t/2$. Observe that $p_k = k!c_k$, where c_k counts these paths with the extra restriction of using the k coordinates in order: the path does not use the second coordinate until after using the first coordinate; the path does not use the third coordinate until after using the second coordinate; etc. We computed c_k for all k simultaneously by a standard pruned combinatorial search. Note, however, that there are faster algorithms, at least for $k = t/2$ (see the a_r formula below).

The number of sequences $(h(R_0), h(R_1), \ldots, h(R_{t-1}))$ producing a cycle of length exactly t is

$$\binom{n}{t/2}p_{t/2} + \cdots + \binom{n}{3}p_3 + \binom{n}{2}p_2 + \binom{n}{1}p_1.$$

Indeed, for each sequence, consider the set S of indices i for which $\pm T_i$ appears in the sequence, and write k for the size of S. Map these T_i, in increasing order of i, to the k basis vectors of $\mathbb{Z}^k$, obtaining a t-step path in $\mathbb{Z}^k$ that starts at $(0, 0, \ldots, 0)$ and that uses all k coordinates. Saying that $R_t = R_0$ is equivalent to saying that the path ends at $(0, 0, \ldots, 0)$, and saying that there is no earlier cycle is equivalent to saying that the path is self-avoiding. Conversely, each such path corresponds to exactly $\binom{n}{k}$ sequences, where $\binom{n}{k}$ accounts for the number of choices of S for this path.

Dividing the above number by $(2n)^t$ gives the probability of a cycle of length exactly t. This probability has the form $(c_{t/2}/2^t)\delta^{t/2} + \cdots + (\cdots)\delta^{t-1}$. This whole computation treats n as a polynomial variable.

The leading coefficient $c_{t/2}/2^t$ has been studied before. (We say "leading" here since $\delta < 1$ and thus the smallest exponent of δ dominates the probability

for large n.) Specifically, define $a_1 = 1$ and $a_r = (r-1)\sum_{k=1}^{r-1} a_k a_{r-k}$ for $r \geq 2$; then a_r is the number of (strongly) "irreducible diagrams with $2r$ nodes", sequence A000699 in the On-line Encyclopedia of Integer Sequences [28], and $c_{t/2} = 2^{t/2} a_{t/2}$. As t grows, the leading coefficient $c_{t/2}/2^t = a_{t/2}/2^{t/2}$ grows much more slowly than the upper bound in [12].

The coefficient c_2 has also been studied before: $p_2 = 2c_2$ is the number of "t-step 2-dimensional closed self-avoiding paths on square lattice", sequence A010566. We are not aware of literature on the intermediate coefficients.