An extended abstract of this paper appears in 7th Algorithmic Number Theory Symposium (ANTS VII), *F. Hess, S. Pauli, and M. Pohst, editors, LNCS 4076, pp. 582–598, 2006. This is the full and corrected version.*

Hard Instances of the Constrained Discrete Logarithm Problem

Ilya Mironov[1] Anton Mityagin[2,4] Kobbi Nissim[3,4]

Abstract

The discrete logarithm problem (DLP) generalizes to the constrained DLP, where the secret exponent x belongs to a set known to the attacker. The complexity of generic algorithms for solving the constrained DLP depends on the choice of the set. Motivated by cryptographic applications, we study sets with succinct representation for which the constrained DLP is hard. We draw on earlier results due to Erdös et al. and Schnorr, develop geometric tools such as generalized Menelaus' theorem for proving lower bounds on the complexity of the constrained DLP, and construct sets with succinct representation with provable non-trivial lower bounds.

1 Introduction

One of the most important assumptions in modern cryptography is the hardness of the discrete logarithm problem (DLP). The scope of this paper is restricted to groups of prime order p, where the DLP is the problem of computing x given (g, g^x) for x chosen uniformly at random from $\mathbb{Z}_p$ (see the next section for notation). In some groups the DLP is believed to have average complexity of $\Theta(\sqrt{p})$ group operations. The *constrained DLP* is defined as the problem of computing x given (g, g^x) where x is chosen uniformly at random from a publicly known set $S \subseteq \mathbb{Z}_p$.

For the standard DLP there is a well-understood dichotomy between generic algorithms, which are oblivious to the underlying group, and group-specific algorithms. By analogy, we distinguish between generic and group-specific algorithms for the constrained DLP. In

[1]Microsoft Research, Silicon Valley Campus.
[2]Department of Computer Science, University of California at San Diego, La Jolla, CA 92037.
[3]Department of Computer Science, Ben Gurion University, Beer-Sheva 84105, Israel.
[4]The work was done in Microsoft Research, Silicon Valley Campus.

this paper we concentrate on the former kind, i.e., generic algorithms. Our main tool for analysis of generic algorithms is the Shoup-Nechaev generic group model [Sho97, Nec94].

The main motivation of our work is the fundamental nature of the problem and the tantalizing gap that exists between lower and upper bounds on the constrained DLP. A trivial generalization of Shoup's proof shows that the DLP constrained to any set $S \subseteq \mathbb{Z}_p$ has generic complexity $\Omega(\sqrt{|S|})$ group operations. On the other hand, Schnorr demonstrates that the DLP constrained to a *random* S of size $\sqrt{p}$ has complexity $\tilde{\Theta}(\sqrt{p}) = \tilde{\Theta}(|S|)$ with high probability [Sch01]. *Explicit* (de-randomized) constructions or even succinct representation of small sets with high complexity, or any complexity better than the square root lower bound were conspicuously absent.

The importance of improving the square root lower bound for concrete subsets of $\mathbb{Z}_p$ is implicit in [Yac98, HS03, SJ04], which suggest exponentiation algorithms that are faster than average for exponents sampled from certain subsets. These algorithms either rely on heuristic assumptions of security of the DLP constrained to their respective sets or use the square root lower bound to the detriment of their efficiency. For example, Yacobi proposes to use "compressible" exponents whose binary representation contains repetitive patterns [Yac98], which can be exploited by some algorithms for fast exponentiation. However, without optimistic assumptions about the complexity of the DLP constrained to this set the method offers no advantage over the sliding window exponentiation. Another method of speeding up exponentiation is to generate an exponent together with a short addition chain for it [Knu97, Ch. 4.6.3]. Absent reliable methods of sampling addition chains with uniformly distributed last elements, this approach depends on the hardness of the DLP on a non-uniform distribution.

The main technical contribution of our work is the proof that the DLP constrained to a set S, which is chosen from an easily sampleable family of sets of cardinality $p^{1/12-\varepsilon}$, has complexity $\Omega(|S|^{3/5})$ with probability $1 - 6p^{-12\varepsilon}$. At a higher level of abstraction we develop combinatorial techniques to bound the complexity of the constrained DLP, which is a global property, using the set's local properties. We view our work as a step towards better understanding the constrained DLP and possibly designing fast exponentiation algorithms tuned to work on exponents from "secure" subsets.

The structure of the paper is as follows. In Section 2 we present a number of results which are known but otherwise scattered in the literature. In Sections 3 and 4 we give new constructions of sets with provable lower bounds on various families of algorithms for solving the constrained DLP.

1.1 Notation

We use the standard notation for asymptotic growth of functions, where

$$
\begin{aligned}
O(g) &= \{f\colon \mathbb{N} \to \mathbb{R}^+ \mid \exists c, n_0 > 0 \text{ s. t. } 0 \le f(n) \le cg(n) \text{ for all } n > n_0\};\\
\Omega(g) &= \{f\colon \mathbb{N} \to \mathbb{R}^+ \mid \exists c, n_0 > 0 \text{ s. t. } cg(n) \le f(n) \text{ for all } n > n_0\};\\
\Theta(g) &= \{f\colon f = O(g) \text{ and } g = O(f)\};\\
\tilde{O}, \tilde{\Omega}, \tilde{\Theta} &\text{ — same as } O, \Omega, \Theta \text{ with logarithmic factors ignored;}\\
\mathbb{Z}_p &\text{ — the field of residues modulo prime } p;\\
x \in_R S &\text{ — } x \text{ chosen uniformly at random from } S.
\end{aligned}
$$

1.2 Previous work

Algorithms for solving number-theoretic problems can be grouped into two main classes: generic attacks, applicable in any group, and specific attacks designed for particular groups. The generic attacks on discrete logarithm include the baby-step giant-step attack [Sha71], Pollard's rho and lambda algorithms [Pol78] as well as their parallelized versions [vOW99, Pol00], surveyed in [Tes01]. The specific attacks have sprouted into a field in their own right, surveyed in [SWD96, Odl00].

A combinatorial view on generic attacks on the DLP was first introduced by Schnorr [Sch01]. He suggested the concept of the generic DL-complexity of a subset $S \subseteq \mathbb{Z}_p$ defined as the minimal number of generic operations required to solve the DLP for any element of $\{g^x \mid x \in S\}$. He showed that the generic DL-complexity of random sets of size $m < \sqrt{p}$ is $m/2 + o(1)$. In part our work is an extension of Schnorr's paper. The combinatorial approach to the DLP was further advanced by [CLS03] which gave a characterization of generic attacks on the entire group of prime order.

Systematically the constrained DLP has been studied for two special cases: Exponents restricted to an interval and exponents with low Hamming weight. Pollard's kangaroo method has complexity proportional to the square root of the size of the interval [Pol00]. The running time of a simple Las Vegas baby-step giant-step attack on low-weight exponents is $O(\sqrt{t}\binom{n/2}{t/2})$, where n is the length and t is the weight of the exponent [Hei93] (for a deterministic version see [Sti02], which credits it to Coppersmith). See [CLP05] for cryptanalysis of a similar scheme in a group of unknown order.

Erdös and D. Newman studied the BSGS-1 complexity (in our notation) and asked for constructions of sets with a high (better than a random subset's) BSGS-1 complexity [EN77].

1.3 Generic algorithms

The generic group model introduced by Shoup and Nechaev [Sho97, Nec94] provides access to a group G via a random injective mapping $\sigma : G \to \Sigma$, which *encodes* group elements. The group operation is implemented as an oracle that on input $\langle \sigma(g), \sigma(h), \alpha, \beta \rangle$ outputs $\sigma(g^\alpha h^\beta)$

(for the sake of notation brevity we roll three group operations, group multiplication, group inversion, and group exponentiation, in one). Wlog, we restrict the arguments of the queries issued by algorithms operating in this model to encodings previously output by the oracle.

The discrete logarithm problem for groups of prime order has a trivial formalization in the generic group model:

Given $p, \sigma(g), \sigma(g^x)$ where g has order p and $x \in_R \mathbb{Z}_p$, find x.

The proof sketch of the theorem below, which is essentially the original one due to Shoup, is reproduced here because it lays the ground for a systematic study of complexity of algorithms in the generic group model.

Theorem 1 ([Sho97]) *Let $\mathcal{A}$ be a probabilistic algorithm and m be the number of queries made by $\mathcal{A}$. $\mathcal{A}$ solves the discrete logarithm problem in a group of prime order p with probability*

$$\Pr[\mathcal{A}(p, \sigma(g), \sigma(g^x)) = x] < \frac{(m+2)^2}{2p} + \frac{1}{p}.$$

The probability space is x, $\mathcal{A}$'s coin tosses, and the random function σ.

Proof [sketch] Instead of letting $\mathcal{A}$ interact with a real oracle, consider the following game played by a simulator. The simulator keeps track of two lists of equal length L_1 and L_2: the list of encodings $\sigma_1, \ldots, \sigma_t \in \Sigma$ and the list of linear polynomials $a_1x+b_1, \ldots, a_tx+b_t \in \mathbb{Z}_p[x]$. Initially L_1 consists of two elements σ_1 and σ_2, which are the two inputs of $\mathcal{A}$, and L_2 consists of 1 and x. When $\mathcal{A}$ issues a query $\langle \sigma_i, \sigma_j, \alpha, \beta \rangle$, the simulator fetches the polynomials $a_ix + b_i$ and $a_jx + b_j$ from L_2, computes $a = \alpha a_i + \beta a_j$ and $b = \alpha b_i + \beta b_j$ and looks up $ax + b$ in L_2. If $ax + b = a_kx + b_k$ for some k, the simulator returns σ_k as the answer to the query. Otherwise, the simulator generates a new element $\sigma_{t+1} \in_R \Sigma \setminus L_1$, appends σ_{t+1} to L_1 and $ax + b$ to L_2, and returns σ_{t+1}.

$\mathcal{A}$ terminates by outputting some $y \in \mathbb{Z}_p$. The game completes as follows:

1. The simulator randomly selects $x^* \in_R \mathbb{Z}_p$.
2. Compute $a_ix^* + b_i$ for all $i \leq m + 2$. If $a_ix^* + b_i = a_jx^* + b_j$ for some $i \neq j$, the simulator fails.
3. $\mathcal{A}$ succeeds if and only if $x^* = y$.

Observe that the game played by the simulator is indistinguishable from the transcript of $\mathcal{A}$'s interaction with the actual oracle unless the simulator fails in step 2. Since for any two distinct polynomials $a_ix + b_i$ and $a_jx + b_j$ the probability that $a_ix^* + b_i = a_jx^* + b_j$ is at most $1/p$, the probability that step 2 fails is at most $(m + 2)^2/2p$. Finally, we observe that the probability that $\mathcal{A}$ wins the game in step 3 is exactly $1/p$, which completes the proof$\square$

It follows from the proof that the probability of success of any probabilistic adaptive algorithm for solving the discrete logarithm in $\mathbb{Z}_p$ in the generic group model can be computed given the list of the linear polynomials induced by its queries. This observation leads us to the concept of generic complexity defined in the next section.

2 Generic Complexity

Definition 1 (Intersection set) *For a set of pairs $L \subseteq \mathbb{Z}_p^2$, we define its* intersection set

$$I(L) = \{x \in \mathbb{Z}_p \mid \exists (a,b),(a',b') \in L \text{ s.t. } ax+b = a'x+b' \text{ and } (a,b) \neq (a',b')\}.$$

The set of pairs from the above definition corresponds to the set of queries asked by the generic algorithm. Its intersection set is the set of inputs on which the simulator from the proof of Theorem 1 fails.

Definition 2 (L recognizes an α-fraction of S) *For $L \subseteq \mathbb{Z}_p^2$, $S \subseteq \mathbb{Z}_p$, and $0 < \alpha \leq 1$ we say that L* recognizes an α-fraction of *the set S if*

$$|S \cap I(L)| \geq \alpha |S|.$$

Definition 3 (Generic complexity) *The set $S \subseteq \mathbb{Z}_p$ is said to have* generic α-complexity *m denoted as $\mathcal{C}_\alpha(S)$ if m is the smallest cardinality of a set L recognizing an α-fraction of S.*

Our definition of generic complexity is slightly different from a similar concept of the generic DL-complexity put forth by Schnorr. We only require that the intersection set $I(L)$ covers a constant fraction of S rather than the entire set [Sch01]. Our definition better matches the standard practice of cryptanalysis, when an attack is considered successful if it succeeds on a nontrivial fraction of the inputs. Moreover, our bounds exhibit different scaling behavior as a function of α, and by parametrizing the definition with α we make the dependency explicit.

Proposition 1 ([Sch01]) *For any $S \subseteq \mathbb{Z}_p$ the generic α-complexity of the set S is bounded as*

$$\sqrt{2\alpha|S|} < \mathcal{C}_\alpha(S) \leq \alpha|S|/2 + 3.$$

Proof The lower bound follows from the fact that for any $L \subseteq \mathbb{Z}_p^2$ the cardinality of the intersection set is bounded as $|I(L)| < |L|^2/2$. Therefore, in order to cover at least an α-fraction of the set, $|L|^2/2$ must be more than $\alpha|S|$.

The upper bound is attained by the following construction. If $2m \geq \alpha|S|$ and $\{x_1, \ldots, x_{2m}\} \subseteq S$, then an α-fraction of S is recognized by L of size $m+2$ defined as

$$L = \left\{(0,0),(0,1),(\frac{1}{x_2 - x_1}, \frac{x_1}{x_1 - x_2}), \ldots, (\frac{1}{x_{2m} - x_{2m-1}}, \frac{x_{2m-1}}{x_{2m-1} - x_{2m}})\right\},$$

since x_i and x_{i-1} are the x-coordinates of the points of intersection of the line $(\frac{1}{x_i - x_{i-1}}, \frac{x_{i-1}}{x_{i-1} - x_i})$ with lines $y = 0$ and $y = 1$ respectively. $\square$

Proposition 2 $\sqrt{2\alpha p} < \mathcal{C}_\alpha(\mathbb{Z}_p) \leq 2\lceil\sqrt{\alpha p}\rceil$.

Proof The lower bound on $\mathcal{C}_\alpha(\mathbb{Z}_p)$ is by Proposition 1. The upper bound is given by the set $L = \{(0,i),(1,-\lambda i) \mid 0 \le i < \lambda\}$, where $\lambda = \lceil\sqrt{\alpha p}\rceil$. Indeed,

$$I(L) = \bigcup_{0\le i,j<\lambda} I(\{(0,i),(1,-\lambda j)\} = \bigcup_{0\le i,j<\lambda} \{\lambda j + i\},$$

which covers $[0,\alpha p)$. $\square$

A tighter (up to a constant factor) bound in the general case and exact values for $\mathcal{C}_1(\mathbb{Z}_p)$ for small primes $p < 100$ appear in [CLS03].

Since the generic complexity is a monotone property, it follows that for any set $S \subseteq \mathbb{Z}_p$

$$\mathcal{C}_\alpha(S) \le \min(\alpha|S|/2 + 3, 2\lceil\sqrt{\alpha p}\rceil).$$

Now we are ready to establish the connection between the generic complexity of a set and the discrete logarithm problem.

Theorem 2 *Let $S \subseteq \mathbb{Z}_p$, $\mathcal{A}_S$ be a generic algorithm that makes $m < \mathcal{C}_\alpha(S)$ queries and outputs a number from $\mathbb{Z}_p$. Then its probability of success is bounded as*[1]

$$\Pr[\mathcal{A}_S(\sigma(g),\sigma(g^x)) = x] < \alpha + \frac{1}{|S|},$$

where the probability is taken over $\mathcal{A}$'s random tape, the oracle answers, and $x \in_R S$. The above bound is tight, i.e., for any set S there is a generic algorithm whose query complexity is $\mathcal{C}_\alpha(S)$ and probability of success is at least $\alpha + 1/|S|$.

Proof [sketch] The proof essentially follows that of Theorem 1. Let L be a set of pairs (a_i, b_i) constructed by the simulator and $x^* \in_R S$ be its choice for x. The adversary succeeds in two cases: either x^* belongs to the intersection set of L or x^* is the output of $\mathcal{A}_S$. The first probability is at most $|I(L) \cap S|/|S| < \alpha$ as long as $m < \mathcal{C}_\alpha(S)$, the second probability is exactly $1/|S|$.

The tightness property follows from the definition of generic complexity. Let L be the set of pairs of size $\mathcal{C}_\alpha(S)$ so that $|S \cap I(L)| \ge \alpha|S|$. Query the oracle $\langle\sigma(g^x),\sigma(g),a,b\rangle$ for all pairs $(a,b) \in L$. With probability $|I(L) \cap S|/|S|$ there is a collision that gives away x, otherwise make a guess that succeeds with probability $1/|S \setminus I(L)|$. $\square$

Notice that the theorem above is unconditional and the adversary is computationally unbounded. In particular, the adversary is given full access to S and can design an S-specific algorithm. As long as the algorithm has only oracle access to the group, $\mathcal{C}_\alpha(S)$ is a lower bound on the number m of oracle queries needed by the algorithm to succeed with probability at least $\alpha + 1/|S|$.

We know that $\mathcal{C}_\alpha(S)$ can be negligible compared to $|S|$ (for instance, according to Proposition 2, when $S = \mathbb{Z}_p$, $|S| = p$ but its generic complexity is $O(\sqrt{p})$). Since the generic complexity is intimately related to the query complexity of any discrete logarithm-solving algorithm, we would like to build sets with higher generic complexity. The next theorem demonstrates that for a fixed p a random set of size less than $\sqrt{p}$ has a near-linear generic complexity.

[1]This statement is stronger than the one in the proceedings version.

Theorem 3 *For a random subset $S \subseteq_R \mathbb{Z}_p$ of size p^ε for some constant $\varepsilon \leq 1/2$ its generic α-complexity is at least*

$$\mathcal{C}_\alpha(S) > \frac{\alpha|S|}{\ln p}$$

with probability $1 - 1/p$ for large enough p.

Proof The proof is by a counting argument. We shall bound the number of the sets S of size $k = p^\varepsilon$ whose α-fraction can be recognized by a set L of size δk, when $\varepsilon \leq 1/2$ and $\delta = \alpha/\ln p$. Suppose $|L| = \delta k$ and $|I(L) \cap S| \geq \alpha k$, where S is to be constructed. There are $\binom{p^2}{\delta k}$ subsets $L \subseteq \mathbb{Z}_p^2$ of size δk. The intersection set $I(L)$ has size at most $(\delta k)^2$ and contains at least αk elements which belong to S. There are thus at most $\binom{p^2}{\delta k}\binom{(\delta k)^2}{\alpha k}$ distinct possibilities for these αk elements. The $(1-\alpha)$-fraction of S can be chosen arbitrarily from $\mathbb{Z}_p$, in $\binom{p}{(1-\alpha)k}$ many ways. In total the number of subsets S of generic complexity δk and cardinality k is bounded by $\binom{p^2}{\delta k}\binom{(\delta k)^2}{\alpha k}\binom{p}{(1-\alpha)k}$. Using $\binom{n}{k} < (ne/k)^k$ for any $0 < k \leq n$ and $x^{-x} < 1.5$ for any $x > 0$ we bound the product as

$$p^{2\delta k+(1-\alpha)k}\delta^{-\delta k+2\alpha k}k^{-\delta k+2\alpha k-k}\alpha^{-\alpha k}(1-\alpha)^{-(1-\alpha)k}e^{\delta k+k} < \left[4p^{2\delta+(1-\alpha)}\delta^{2\alpha}k^{-\delta+2\alpha-1}e^{\delta+1}\right]^k < \left[12p^{2\delta+(1-\alpha)}\delta^{2\alpha}k^{-\delta+2\alpha-1}\right]^k.$$

We want this number to be less than a $1/p$-fraction of the number of subsets of $\mathbb{Z}_p$ of size k, which is $1/p\binom{p}{k} > 1/p(p/k)^k$. By taking the kth root of both numbers and substituting $k = p^\varepsilon$, we arrive at the following inequality:

$$12\delta^{2\alpha}p^{2\delta+(1-\alpha)+\varepsilon(-\delta+2\alpha-1)+p^{-\varepsilon}} < p^{1-\varepsilon}.$$

Notice that the inequality holds for $\delta < \alpha(1-2\varepsilon)$ if $\varepsilon < 1/2$ and for $\delta < \alpha/\ln p$ if $\varepsilon = 1/2$. When ε is constant and p is large enough, $\delta = \alpha/\ln p < \alpha(1-2\varepsilon)$. □

The bottom line of the theorem we just proved is that hard sets (where the discrete logarithm is hard to compute using a generic algorithm) are easy to come by. In fact, almost any set has high generic complexity (also previously observed in [Sch01]).

In what follows we sharply lower the amount of randomness that is required to provide any non-trivial guarantee of generic complexity.

3 More complexities and lower bounds

Many sets of group elements with special properties may be attacked using a baby-step giant-step method. In this method the attacker first computes $g^{c_1},\ldots,g^{c_m}$ (giant steps) and then compares them against $g^{a_1x+b_1},\ldots,g^{a_mx+b_m}$ (baby steps). Any collision between a baby step and a giant step gives away x. We define the complexity of this method along the lines of the generic complexity from the previous section.

Definition 4 (Intersection set-2) *For a set of pairs $L \subseteq \mathbb{Z}_p^2$ and a set of points $C \subseteq \mathbb{Z}_p$, we define their* intersection set *as*

$$I(L, C) = \{x \in \mathbb{Z}_p \mid \exists (a, b) \in L, c \in C \text{ s.t. } a \neq 0 \text{ and } ax + b = c\}.$$

Definition 5 (Baby-step giant-step complexity.) *The set $S \subseteq \mathbb{Z}_p$ is said to have the* baby-step giant-step α-complexity *(BSGS complexity for short) m denoted as $\mathcal{C}_\alpha^{\text{bsgs}}(S)$ if m is the smallest integer such that there exist $L \subseteq \mathbb{Z}_p^2$ and $C \subseteq \mathbb{Z}_p$, with $|L| = |C| = m$ and $|I(L, C) \cap S| \geq \alpha|S|$.*

An important particular case of the baby-step giant-step method is when all lines defined by L are parallel (i.e., all $a_i = 1$).

Definition 6 (BSGS-1 complexity) *The set $S \subseteq \mathbb{Z}_p$ has* BSGS-1 α-complexity *m denoted as $\mathcal{C}_\alpha^{\text{bsgs1}}(S)$ if m is the smallest integer such that there exist $L \subseteq \{1\} \times \mathbb{Z}_p$ and $C \subseteq \mathbb{Z}_p$, with $|L| = |C| = m$ and $|I(L, C) \cap S| \geq \alpha|S|$.*

Equivalently, $\mathcal{C}_\alpha^{\text{bsgs1}}(S)$ is the smallest integer n such that there exist $X, Y \subseteq \mathbb{Z}_p$, with $n = |X| = |Y|$ and $|S \cap (X - Y)| \geq \alpha|S|$, where $X - Y$ is the set of pairwise differences between X and Y. The intersection sets from the three definitions of complexities appear in Fig. 1.

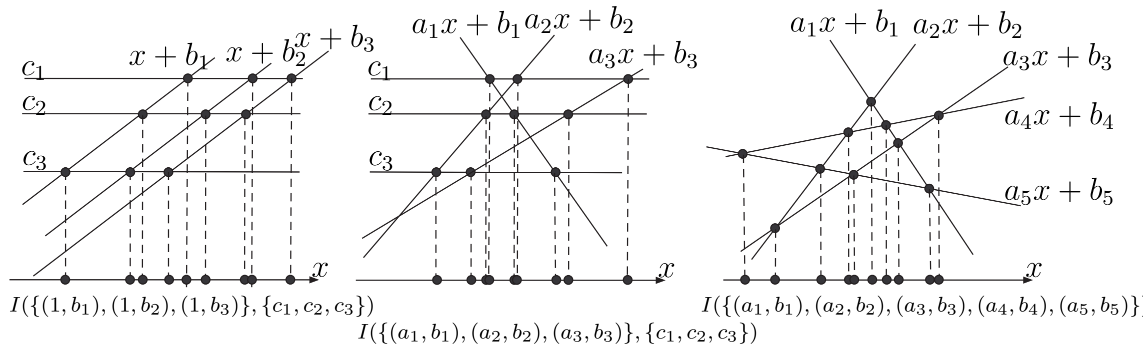

Figure 1: Intersection sets for BSGS-1, BSGS, and generic complexities.

The problem of computing $\mathcal{C}_\alpha^{\text{bsgs1}}(S)$ is superficially similar to a number of problems in additive number theory concerned with studying properties of $X - Y$. However, our goal is fundamentally different since we require that $X - Y$ cover a large fraction of S rather than be its exact equal. To the best of our knowledge, the only paper in the literature directly applicable to bounding $\mathcal{C}_1^{\text{bsgs1}}(S)$ is a 1977 paper by Erdös and Newman [EN77]. They proved analogues of our Theorems 3 and 6 and bounded the BSGS-1 complexity (in our notation) of the set of small squares $\{x^2 \mid x < \sqrt{p}\}$ to be $\Omega(p^{1/3 - c/\log\log p})$. They leave open the problem of constructing sets with a strictly linear BSGS-1 complexity (without the $1/\log p$ factor).

The BSGS and BSGS-1 complexities provide useful upper bounds for the generic complexity.

Proposition 3 $\frac{1}{2}\mathcal{C}_\alpha(S) \le \mathcal{C}_\alpha^{\text{bsgs}}(S) \le \mathcal{C}_\alpha^{\text{bsgs1}}(S)$.

Proof Let $C' = \{0\} \times C = \{(0, c) \mid c \in C\}$. Then $I(L, C) \subseteq I(L \cup C')$, which implies the first inequality. The second inequality follows from the fact that any BSGS-1 attack is also a BSGS attack. □

Consider, for example, the baby-step giant-step attacks on exponents with low Hamming weight [Hei93, Sti02]. Define $S_\lambda = \{x \in \mathbb{Z}_p \mid \nu(x) = \lambda|x|\}$, where $\nu(x)$ is the number of ones in the binary representation of x. Stinson [Sti02] estimates the complexity of the randomized algorithm due to Coppersmith to yield

$$\mathcal{C}_{1/2}^{\text{bsgs1}}(S_\lambda) = \tilde{O}(p^{1/2 \log_2(\lambda^{-\lambda}(1-\lambda)^{\lambda-1})}).$$

For instance, if $\lambda = 1/4$, the bound becomes $\mathcal{C}_{1/2}^{\text{bsgs1}}(S_{0.25}) = \tilde{O}(p^{0.406})$.

Following Propositions 2, 3, and Theorem 3 the BSGS-1 α-complexity of a set of cardinality less than $\sqrt{p}$ lies between $\sqrt{\alpha|S|/2}$ and $2\sqrt{\alpha p}$, where the lower bound is trivial and the upper bound is approximated up to a logarithmic factor by almost any subset of size $\sqrt{p}$. In this section we construct a set with succinct representation and a non-trivial BSGS-1 complexity. We start by stating without proof an important combinatorial lemma known as the Zarankiewicz problem [Zar51]:

Theorem 4 *[Bol98, Ch. IV.2] Let $Z(n, s, t)$ be the maximum number of ones that can be arranged in an $n \times n$ matrix such that there is no all-one $t \times s$ (possibly disjoint) submatrix. Then*

$$Z(n, s, t) < s^{1/t} n^{2-1/t}.$$

Notice that the asymptotic of the bound on $Z(n, s, t)$ depends on the smallest of the two dimensions of the prohibited all-one submatrix. It is known that the bound is tight (up to a constant factor) for $t = 2, 3$.

Our second combinatorial tool follows from a more general upper bound due to A. Naor and Verstraëte on the number of edges in a bipartite graph without cycles of length $2k$ (C_{2k}-free graph):

Theorem 5 ([NV05]) *The maximum number of edges in a C_{2k}-free (n, n)-bipartite graph is less than $2kn^{1+1/k}$.*

When $k = 2$ the two theorems overlap. Indeed, a 0-1 matrix is also a bipartite graph, where the rows and columns form the vertex set and the non-zero elements indicate adjacency of corresponding vertices. In this case an all-one 2×2 submatrix represents a cycle of length 4 in the graph. Our theorems fully reflect this relationship: Theorem 6 can be proved using either the Zarankiewicz or the Naor-Verstraëte bound; its generalization Theorem 7 makes use of C_{2k}-free graphs, while Theorems 8 and 9 apply the Zarankiewicz bound.

Theorem 6 *Suppose $S \subseteq \mathbb{Z}_p$ has the property that all pairwise sums of different elements of S are distinct. Then*

$$\mathcal{C}_\alpha^{\text{bsgs1}}(S) > (\alpha|S|/\sqrt{2})^{2/3}.$$

Proof Take $X, Y \subseteq \mathbb{Z}_q$, such that $n = |X| = |Y|$ and $|S \cap (X - Y)| > \alpha|S|$. Consider an $n \times n$ matrix M, whose rows and columns are labeled with elements of X and Y respectively. For each element $s \in S \cap (X - Y)$ find one pair $x \in X$ and $y \in Y$ such that $s = x - y$ and set the (x, y) entry of the matrix to one. Since $X - Y$ covers at least an α-fraction of S, the number of ones in the matrix is at least $\alpha|S|$.

We claim that M does not contain an all-one 2×2 submatrix. Assume the opposite: The submatrix given by elements x_1, x_2 and y_1, y_2 has four ones. It follows that all four $s_{11} = x_1 - y_1$, $s_{12} = x_1 - y_2$, $s_{21} = x_2 - y_1$, $s_{22} = x_2 - y_2 \in S$. Then $s_{11} + s_{22} = s_{12} + s_{21}$, which contradicts the assumption that all pairwise sums of elements of S are distinct. Applying the Zarankiewicz bound for the case $s = t = 2$, we prove that

$$\alpha|S| < Z(n, 2, 2) < \sqrt{2}n^{2-1/2} = \sqrt{2}n^{3/2},$$

which implies that $n = \mathcal{C}_\alpha^{\text{bsgs1}}(S) > (\alpha|S|/\sqrt{2})^{2/3}$. □

The sets where sums of pairs of different elements are distinct are known in combinatorics as weak Sidon sets. They are closely related to (strong) Sidon sets, also called B_2 sequences, where all pairwise sums (of not necessarily different elements) are distinct (for a comprehensive survey see [O'B04] that includes more than 120 bibliographic entries). Explicit constructions of Sidon subsets of $\{1, \ldots, n\}$ due to Singer and Ruzsa have cardinality at least $\sqrt{n} - n^{.263}$ [Sin38, BC63, Ruz93, BHP01].

We additionally require that the sums be different modulo p. The size of such sets is bounded from above by $p^{1/2} + 1$ [HHÖ04, Theorem 3]. The easiest shortcut to constructing weak modular Sidon sets is to take a strong Sidon subset of $\{0 \ldots \lfloor p/2 \rfloor\}$ (see also [O'B02, Ch. 3] and [GS80]). Denser Sidon sets may be constructed for primes of the form $p = q^2 + q + 1$, where q is also prime [Sin38]. Existence of infinitely many such primes is implied by Schinzel's Hypothesis H and their density follows from the even stronger Bateman-Horn conjecture [Guy04, A]. Interestingly, modular Sidon sets are useful not only in constructing sets with high complexity, via Theorem 6, but also for solving the discrete logarithm problem in $\mathbb{Z}_p$ [CLS03].

4 Beyond the Basics

Theorem 6 can be generalized to make use of Sidon sets of higher order. First, we prove that if all k-wise sums of elements of S are distinct (counting permutations of the same k-tuple only once), then there is a bound on the BSGS-1 complexity. Second, we provide a result that there exist such sets of size $\Theta(p^{1/k})$.

Theorem 7 *Suppose $S \subseteq \mathbb{Z}_p$ is such that all k-wise sums of different elements of S are distinct (excluding permutation of the summands). Then*

$$\mathcal{C}_\alpha^{\text{bsgs1}}(S) > (\alpha|S|/(2k))^{k/(k+1)}.$$

Proof Take $X, Y \subseteq \mathbb{Z}_p$, such that $\mathcal{C}_\alpha^{\text{bsgs1}}(S) = |X| = |Y| = n$ and $|S \cap (X - Y)| > \alpha|S|$. Instead of the matrix as in Theorem 6, consider a bipartite graph $G(X, Y)$, where there is an edge (x, y) if and only if $x - y \in S$ (keep only one edge per element of S).

We claim that there are no $2k$-cycles (without repetitive edges) in the bipartite graph G. Assume the opposite: There is a cycle $(x_1, y_1, \ldots, x_k, y_k, x_1, y_1)$. Consider two sums: $(x_1 - y_1) + (x_2 - y_2) + \cdots + (x_k - y_k)$ and $(x_2 - y_1) + (x_3 - y_2) + \cdots + (x_k - y_{k-1}) + (x_1 - y_k)$. Not only are the two sums equal, they also consist of k elements of S each, and these elements are all distinct (as every element of S appears as an edge of G at most once). A contradiction is found.

The number of edges in an (n, n)-bipartite graph without $2k$-cycles is less than $2kn^{1+1/k}$ (Theorem 5). Therefore $\alpha|S| < 2kn^{1+1/k}$, and $n = \mathcal{C}_\alpha^{\text{bsgs1}}(S) > (\alpha|S|/(2k))^{k/(k+1)}$. □

Bose and Chowla give a construction for subsets of $\{1, \ldots, q^k\}$ of prime size q whose k-wise sums are distinct (in integers, not modulo p) [BC63]. By choosing the largest prime q less than $p^{1/k}$ (which, for large p is more than $p^{1/k} - p^{0.525/k}$ [BHP01]) an interval of length q^k/k with a $1/k$ proportion of the set's elements, we guarantee that all k-sums are distinct in $\mathbb{Z}_p$ as well. Unfortunately, [BC63] does not provide an efficient sampling algorithm.

Along the lines of Theorem 6 we prove that other verifiable criteria imply non-trivial bounds on the BSGS and generic complexity.

Theorem 8 *Suppose $S \subseteq \mathbb{Z}_p$ is such that for any distinct $x_1, x_2, y_1, y_2, z_1, z_2 \in S$:*

$$\det \begin{pmatrix} x_1 - y_1 & x_2 - y_2 \\ y_1 - z_1 & y_2 - z_2 \end{pmatrix} \neq 0. \tag{1}$$

Then

$$\mathcal{C}_\alpha^{\text{bsgs}}(S) > (\alpha|S|/\sqrt{3})^{2/3}.$$

Proof Take $L \subseteq \mathbb{Z}_p^2$ and $C \subseteq \mathbb{Z}_p$, such that $|L| = |C| = n$ and $|I(L, C) \cap S| > \alpha|S|$. As in Theorem 6 consider the $n \times n$ matrix M whose rows and columns are labeled with elements of L and C respectively. For each element $s \in S \cap I(L, C)$ set one entry in row (a, b) and column c to one, where $s = (c - b)/a$. Thus, the total number of ones in the matrix is exactly $m = |I(L, C) \cap S|$. If there is a 2×3 all-one submatrix in M, then property (1) does not hold (three parallel lines divide two other lines proportionally). The Zarankiewicz bound implies that

$$\alpha|S| < Z(n, 3, 2) < \sqrt{3}n^{2-1/2} = \sqrt{3}n^{3/2}.$$

Hence $\mathcal{C}_\alpha^{\text{bsgs}}(S) = n \geq (\alpha|S|/\sqrt{3})^{2/3}$. □

Constructing a large subset of $\mathbb{Z}_p$ with short description satisfying the condition of the previous theorem is a difficult problem. Fortunately, the probability that a random 6-tuple of $\mathbb{Z}_p$ elements fails to satisfy (1) is $2/p$ [Sch80]. This observation motivates the following definition:

Definition 7 ($\mathcal{S}(N, k)$ family) *Let $\mathcal{S}(N, k) = \{x_1, \ldots, x_N\}$ be a family of subsets of $\mathbb{Z}_p$, where $x_1, \ldots, x_N : \mathcal{K} \mapsto \mathbb{Z}_p$ are k-wise independent random variables ($\mathcal{K}$ is the probability space).*

Properties of $\mathcal{S}(N,k)$ are established in the following proposition:

Proposition 4 *1. $\mathcal{S}(N,k)$ can be defined over $\mathcal{K} = \mathbb{Z}_p^k$.*

2. For $k > 1$, $\Pr_{S \in_R \mathcal{S}(N,k)}[|S| \neq N] < N^2/p$.

3. If $h \in \mathbb{Z}[y_1, \ldots, y_k]$ and $d = \deg(h) > 0$, then

$$\Pr_{S \in \mathcal{S}(N,k)}[\exists \text{ distinct } z_1, \ldots, z_k \in S \text{ with } h(z_1, \ldots, z_k) = 0] < N^k d/p.$$

Proof 1. To construct $\mathcal{S}(N,k)$ we use a well-known k-universal family of functions (following [CW77]). Let the probability space be $\mathcal{K} = \mathbb{Z}_p^k$ and $f_a(x) = a_{k-1}x^{k-1} + \cdots + a_0$ for $a = (a_0, \ldots, a_{k-1}) \in \mathcal{K}$. Define the random variables $x_i = f_a(i) \colon \mathcal{K} \to \mathbb{Z}_p$ for $1 \leq i \leq N$. We claim that the variables $x_1, \ldots, x_N$ are k-wise independent. This follows from the system $f_a(i_1) = y_1, \ldots, f_a(i_k) = y_k$ having a unique solution $a \in \mathcal{K}$ for any distinct $i_1, \ldots, i_k \in \{1, \ldots, N\}$ and $y_1, \ldots, y_k \in \mathbb{Z}_p$. Notice that any $S \in \mathcal{S}(N,k)$ can be easily enumerated and sampled from.

2. Let I_{ij} be the indicator variable, which is equal to 1 when $x_i = x_j$ and 0 otherwise. The cardinality of $S = \{x_1, \ldots, x_N\}$ is at least $N - \sum_{i<j} I_{ij}$. Since x_i and x_j are independent for all $i \neq j$, $E[I_{ij}] = 1/p$. By linearity of expectation, the expected value $E[\sum_{i<j} I_{ij}] < N^2/p$. By Markov's inequality $\Pr[|S| \neq N] = \Pr[\sum_{i<j} I_{ij} \geq 1] < N^2/p$.

3. Let $I_{i_1,\ldots,i_k}$ for all distinct $1 \leq i_1, \ldots, i_k \leq N$ be the indicator variable that is 1 if and only if $h(x_{i_1}, \ldots, x_{i_k}) = 0$. By independence of the variables and [Sch80] $E[I_{i_1,\ldots,i_k}] \leq 2/p$, which by linearity of expectation and Markov's inequality implies that $\Pr_S[\exists \text{ distinct } x_1, \ldots, x_k \in S \text{ with } h(x_1, \ldots, x_k) = 0] \leq \Pr_S[\sum_{i_1,\ldots,i_k} I_{i_1,\ldots,i_k} \geq 1] < N^k d/p$. $\square$

It follows that a randomly chosen set from the family $\mathcal{S}(p^{1/6-\varepsilon}, 6)$ has size $p^{1/6-\varepsilon}$ with probability at least $1 - p^{-2/3}$ and satisfies the condition of Theorem 8 with probability at least $1 - 2p^{-6\varepsilon}$.

To apply a similar argument to the all-powerful generic complexity, we may show that for small constants m_1 and m_2, the projections on the x axis of the intersection points of an irregular m_1 by m_2 grid (in which lines need not be parallel) satisfy a certain relationship. Next, a set S, where any m_1m_2-tuple avoids this relationship, is to be constructed.

Let us see first why this argument works for some values of m_1 and m_2, and then improve the parameters. Let $m_1 = 4$ and $m_2 = 5$. There are 9 lines that can be described using 18 parameters. On the other hand, there are 20 points that form the intersection set of these lines. Each of the 20 intersection points imposes a linear equation on the parameters, and hence the system is overdetermined (even if we exclude linearly dependent equations). In particular, this implies that the probability that a random 20-tuple of elements of $\mathbb{Z}_p$ is coverable by a 4×5 grid is negligible. We refine this argument in the following proposition.

Proposition 5 (Bipartite Menelaus' theorem) *Consider seven lines $l_{x,y,z}, l_{1,2,3,4}$ forming an irregular grid, and their twelve intersection points. Let x_i,y_i,z_i be projections on the*

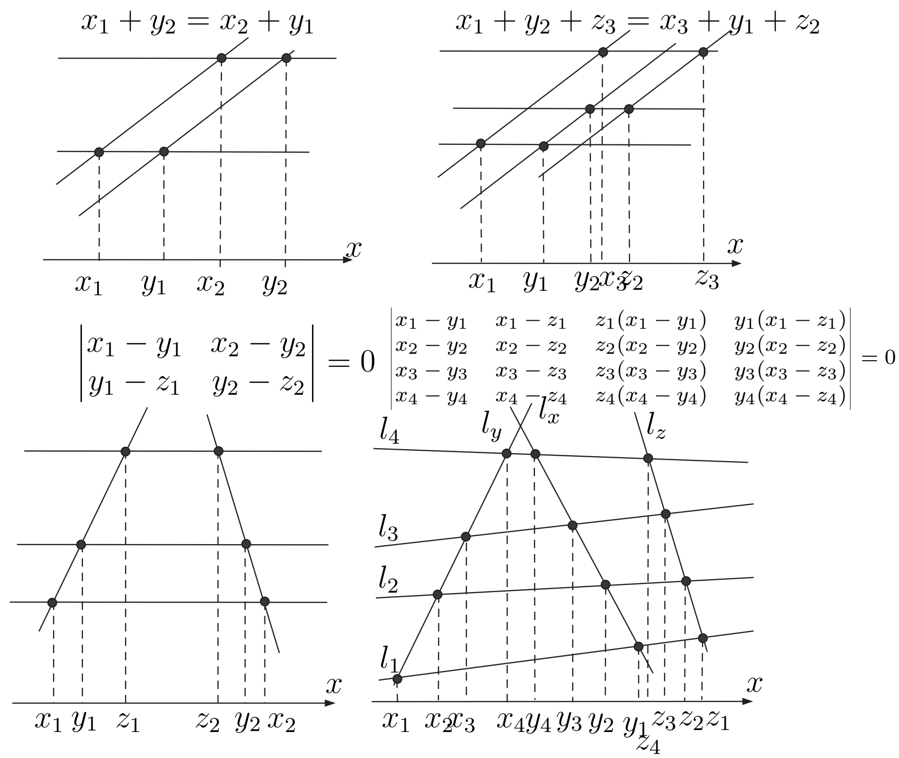

Figure 2: “Prohibited” configurations (Theorems 6, 7 and 8, Proposition 5).

x axis of the intersection points of l_i with lines l_x, l_y, l_z. Then the following holds:

$$\det \begin{pmatrix} x_1 - y_1 & x_1 - z_1 & z_1(x_1 - y_1) & y_1(x_1 - z_1) \\ x_2 - y_2 & x_2 - z_2 & z_2(x_2 - y_2) & y_2(x_2 - z_2) \\ x_3 - y_3 & x_3 - z_3 & z_3(x_3 - y_3) & y_3(x_3 - z_3) \\ x_4 - y_4 & x_4 - z_4 & z_4(x_4 - y_4) & y_4(x_4 - z_4) \end{pmatrix} = 0. \tag{2}$$

Proof Denote the 4×4 matrix in (2) by M. Observe that if any of the seven lines is vertical, (2) follows immediately. Indeed, if $l_y = \{x = \text{const}\}$, then $y_1 = y_2 = y_3 = y_4$ and the second and the fourth columns of M are linearly dependent. Moreover, $\det M$ is invariant under permutations of l_x,l_y, and l_z, which takes care of vertical l_y or l_z. If l_i is vertical for some $1 \le i \le 4$, then the ith row of M is all-zero, and $\det M = 0$.

If none of the lines is vertical, we can write down equations for all of them in the Cartesian coordinates. Let $l_{x,y,z} = \{a_{x,y,z}x + b_{x,y,z}\}$ and $l_i = \{c_i x + d_i\}$ for $1 \le i \le 4$. Each intersection point imposes an equation on the parameters of the two lines incident with it, a total of 12 equations in 14 unknowns. However, the system always has a trivial solution, when all lines are equal. Rewrite the system using new variables: $\tilde{a}_y = a_y - a_x$, $\tilde{b}_y = b_y - b_x$,

$\tilde{a}_z = a_z - a_x$, $\tilde{b}_z = b_z - b_x$, $\tilde{c}_1 = c_1 - a_x$, $\tilde{d}_1 = d_1 - b_x$, etc. The result is a homogenous system of 12 linear equations in 12 new variables. It has a non-zero solution if and only if its matrix is singular (only non-zero elements are shown):

$$M' = \begin{pmatrix} \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & -x_1 & -1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & -y_1 & -1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & -z_1 & -1 \\ x_2 & 1 & \cdot & \cdot & \cdot & \cdot & -x_2 & -1 & \cdot & \cdot & \cdot & \cdot \\ y_2 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & -y_2 & -1 & \cdot & \cdot \\ z_2 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & -z_2 & -1 \\ \cdot & \cdot & x_3 & 1 & \cdot & \cdot & -x_3 & -1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & y_3 & 1 & \cdot & \cdot & \cdot & \cdot & -y_3 & -1 & \cdot & \cdot \\ \cdot & \cdot & z_3 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & -z_3 & -1 \\ \cdot & \cdot & \cdot & \cdot & x_4 & 1 & -x_4 & -1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & y_4 & 1 & \cdot & \cdot & -y_4 & -1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & z_4 & 1 & \cdot & \cdot & \cdot & \cdot & -z_4 & -1 \end{pmatrix}.$$

One can verify that $\det(M) = \det(M')$. □

In the full version of the paper we give a geometric proof of the proposition, deriving (2) directly, and explain the connection with classic Menelaus' theorem. We also give an alternative statement of the theorem, which puts it in the realm of projective geometry.

Proposition 5 is the "minimal" condition that holds for the x-coordinates of the intersection points of two sets of lines in general position. Indeed, it follows from the proof that for any assignment of distinct values to the eleven variables $x_{1,2,3,4}$, $y_{1,2,3,4}$, $z_{1,2,3}$ there is a configuration of lines whose intersection points project to those variables. Other configurations with as many or fewer intersection points do not produce any conditions either. For example, six lines intersecting two lines can project to any collection of twelve points.

All geometric arguments (Theorems 6, 7, and 8, Proposition 5) are illustrated in Fig. 2.

Theorem 9 *If S is chosen from $\mathcal{S}(p^{1/12-\varepsilon}, 12)$, then with probability at least $1 - 6p^{-12\varepsilon}$*

$$\mathcal{C}_\alpha(S) > (\alpha|S|/\sqrt[3]{4})^{3/5}.$$

Proof Consider the set of lines $L \subseteq \mathbb{Z}_p^2$ such that $\mathcal{C}_\alpha(S) = |I(L) \cap S|$ and $n = |L|$. As in Theorem 8, we apply the Zarankiewicz bound to the $n \times n$ matrix, only now both the rows and the columns are labeled with elements of the set L. Similarly, only one occurrence of an element of S as the x-coordinate of the intersection of two distinct lines is recorded in the matrix.

According to Proposition 4, S avoids solutions to the equation (2), whose left-hand side is a multivariate polynomial of total degree 6, with probability greater than $1 - 6p^{-12\varepsilon}$. Therefore the probability that there exist 12 points in S that can be the intersection set of two groups of lines consisting of 3 and 4 lines respectively is less than $6p^{-12\varepsilon}$. Finally, as before, $\alpha|S| < Z(n, 4, 3) < 4^{1/3} n^{2-1/3} = \sqrt[3]{4} n^{5/3} = \sqrt[3]{4}\mathcal{C}_\alpha(S)^{5/3}$. □

Unlike the proofs of Theorems 6, 7, and 8, where the classes in which the lines are grouped arise naturally, the use of bipartite Menelaus' theorem in the analysis of generic complexity above might appear less motivated. In fact, classic Menelaus' theorem imposes a simple condition (a cubic equation) on the intersection set of four lines. It is the second step of the argument, where we translate absence of a certain submatrix (subgraph) into sparseness of the entire matrix, which becomes problematic: Unless H is bipartite, H-free graphs on n vertices may have as many as $\Theta(n^2)$ edges according to the celebrated Turán theorem [Bol98, Ch. IV.2].

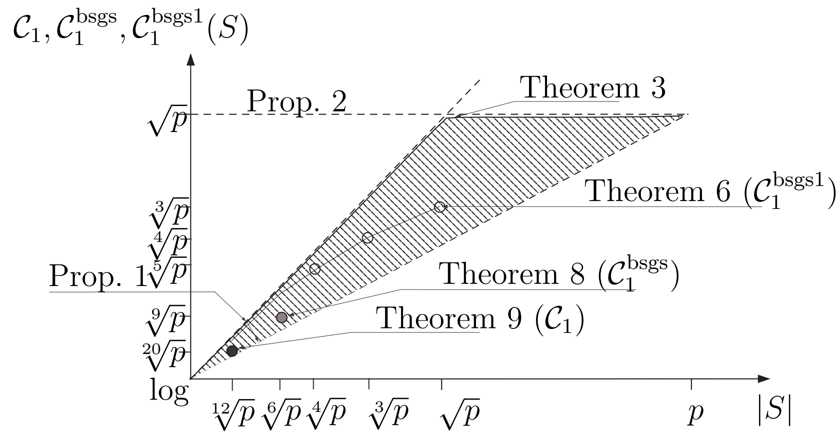

Figure 3: Generic complexities and bounds (in logscale). Propositions 2 and 1 bound the triangle that contains all possible values for generic complexity. The theorems point to lower bounds provable for complexities of their respective constructions.

5 Conclusion

In this paper we develop a theory of lower bounds in the generic group model on the discrete logarithm problem constrained to a subset $S \subseteq \mathbb{Z}_p$ known to the attacker (constrained DLP). We give a first construction of a set with succinct description whose generic complexity is more than the square root of its size (Theorem 9). There exists an apparent gap between our construction ($|S| = p^{1/12}$ and $\mathcal{C}_1(S) = |S|^{3/5}$) and a random set of size $p^{1/2}$ whose complexity is almost linear in its size. Bridging this gap constitutes an interesting open problem whose solution would shed some light on the intrinsic difficulty of the discrete logarithm problem. We also define restricted versions of the generic complexity that capture the complexity of baby-step-giant-step algorithms. We give an explicit, deterministic construction of a collection of sets, whose complexity in respect to the weakest family of baby-step-giant-step algorithms becomes near-optimal as their size decreases (Theorem 7). Various bounds and constructions are put together in Fig. 3.

Acknowledgments. The authors thank Constantin Shramov for his crucial contribution to the geometric proof of Theorem 11, and Fan Chung, Kevin O'Bryant, Imre Ruzsa, and the anonymous reviewer of ANTS VII for their advice and helpful comments.

References

[BC63] Raj C. Bose and Sarvadaman Chowla. Theorems in the additive theory of numbers. *Comment. Math. Helv.*, 37:141–147, 1962–1963.

[BHP01] R. C. Baker, G. Harman, and J. Pintz. The difference between consecutive primes. II. *Proc. London Math. Soc. (3)*, 83(3):532–562, 2001.

[Bol98] Béla Bollobás. *Modern graph theory*, volume 184 of *Graduate texts in mathematics*. Springer, 1998.

[CLP05] Jean-Sébastien Coron, David Lefranc, and Guillaume Poupard. A new baby-step giant-step algorithm and some applications to cryptanalysis. In Josyula R. Rao and Berk Sunar, editors, *CHES 2005*, volume 3659 of *Lecture Notes in Computer Science*, pages 47–60. Springer, 2005.

[CLS03] M. Chateauneuf, Alan Ling, and Douglas R. Stinson. Slope packings and coverings, and generic algorithms for the discrete logarithm problem. *J. Comb. Designs*, 11(1):36–50, 2003.

[CW77] Larry Carter and Mark N. Wegman. Universal classes of hash functions. In *STOC 1977*, pages 106–112, 1977.

[EN77] Paul Erdös and Donald J. Newman. Bases for sets of integers. *J. Number Theory*, 9(4):420–425, 1977.

[GS80] Ronald L. Graham and Neal J.A. Sloane. On additive bases and harmonious graphs. *SIAM J. Algebraic and Discrete Methods*, 1:382–404, 1980.

[Guy04] Richard K. Guy. *Unsolved Problems in Number Theory*. Springer-Verlag, third edition, 2004.

[Hei93] Rafi Heiman. A note on discrete logorithms with special structure. In Rainer A. Rueppel, editor, *Advances in Cryptology—EUROCRYPT '92*, volume 658 of *Lecture Notes in Computer Science*, pages 454–457. Springer, 1993.

[HHÖ04] Harri Haanpää, Antti Huima, and Patric R. J. Östergård. Sets in $\mathbb{Z}_n$ with distinct sums of pairs. *Discrete Applied Mathematics*, 138(1–2):99–106, 2004.

[HS03] Jeffrey Hoffstein and Joseph H. Silverman. Random small Hamming weight products with applications to cryptography. *Discrete Applied Mathematics*, 130(1):37–49, 2003.

[Knu97] Donald E. Knuth. *Seminumerical Algorithms*, volume 2 of *The Art of Computer Programming*. Addison-Wesley, third edition, 1997.

[Nec94] Vassiliy I. Nechaev. Complexity of a determinate algorithm for the discrete logarithm. *Math. Notes*, 55(2):165–172, 1994.

[NV05] Assaf Naor and Jacques Verstraëte. A note on bipartite graphs without $2k$-cycles. *Probability, Combinatorics and Computing*, 14(5–6):845–849, 2005.

[O'B02] Kevin O'Bryant. *Sidon Sets and Beatty Sequences*. PhD thesis, U. of Illinois in Urbana-Champaign, 2002.

[O'B04] Kevin O'Bryant. A complete annotated bibliography of work related to Sidon sequences. *Electr. J. Combinatorics*, DS11, July 2004.

[Odl00] Andrew M. Odlyzko. Discrete logarithms: The past and the future. *Des. Codes Cryptography*, 19(2/3):129–145, 2000.

[Pol78] John M. Pollard. Monte Carlo methods for index computation (mod p). *Mathematics of Computation*, 32:918–924, 1978.

[Pol00] John M. Pollard. Kangaroos, monopoly and discrete logarithms. *J. Cryptology*, 13(4):437–447, 2000.

[Ruz93] Imre Z. Ruzsa. Solving a linear equation in a set of integers. Part I. *Acta Arith.*, 65:259–282, 1993.

[Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, 1980.

[Sch01] Claus-Peter Schnorr. Small generic hardcore subsets for the discrete logarithm. *Inf. Process. Lett.*, 79(2):93–98, 2001.

[Sha71] Daniel Shanks. Class number, a theory of factorization, and genera. In Donald J. Lewis, editor, *1969 Number Theory Institute*, volume 20 of *Proceedings of Symposia in Pure Mathematics*, pages 415–440, Providence, Rhode Island, 1971. American Mathematical Society.

[Sho97] Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *Advances in Cryptology—EUROCRYPT '97*, volume 1233 of *Lecture Notes in Computer Science*, pages 256–266. Springer, 1997.

[Sin38] James Singer. A theorem in finite projective geometry and some applications to number theory. *Trans. Amer. Math. Soc.*, 43:377–385, 1938.

[SJ04] Yaron Sella and Markus Jakobsson. Constrained and constant ratio hash functions. Manuscript, 2004.

[Sti02] Douglas R. Stinson. Some baby-step giant-step algorithms for the low Hamming weight discrete logarithm problem. *Math. Comput.*, 71(237):379–391, 2002.

[SWD96] Oliver Schirokauer, Damian Weber, and Thomas F. Denny. Discrete logarithms: The effectiveness of the index calculus method. In H. Cohen, editor, *ANTS-II*, volume 1122 of *Lecture Notes in Computer Science*, pages 337–361. Springer, 1996.

[Tes01] Edlyn Teske. Square-root algorithms for the discrete logarithm problem (a survey). In *Public-Key Cryptography and Computational Number Theory*, pages 283–301, 2001.

[vOW99] Paul C. van Oorschot and Michael J. Wiener. Parallel collision search with cryptanalytic applications. *J. Cryptology*, 12(1):1–28, 1999.

[Yac98] Yacov Yacobi. Fast exponentiation using data compression. *SIAM J. Comput.*, 28(2):700–703, 1998.

[Zar51] Kazimierz Zarankiewicz. Problem P 101. *Colloq. Math.*, 2:301, 1951.

A Bipartite Menelaus' Theorem

Let us first recap the classic Menelaus theorem.

Theorem 10 (Classic Menelaus) *Consider four directed lines intersecting at six points (see Fig. 4a). Then*

$$AD \cdot BF \cdot CE = BD \cdot CF \cdot AE, \tag{3}$$

where the segments' lengths are signed, i.e., positive if their direction agrees with that of the line they belong to and negative otherwise.

Less known is that (3) is equivalent to the following:

$$\det \begin{pmatrix} DB & AD & AB \\ BC & FC & BF \\ ED & DF & FE \end{pmatrix} = 0. \tag{4}$$

One way to interpret the theorem is to add an x-axis to the drawing and consider projections of the intersection points onto this axis. Since the ratios of signed collinear segments are invariant under orthogonal projection, (3) implies that

$$(x_A - x_D) \cdot (x_B - x_F) \cdot (x_C - x_E) = (x_B - x_D) \cdot (x_C - x_F) \cdot (x_A - x_E), \tag{5}$$

where x_A is the x-coordinate of A, etc.

We may reverse the Menelaus theorem and ask whether a given six-tuple can be the projection of the intersection points of four distinct lines. It is easy to check that (5) is not only necessary but is also a sufficient condition for such four lines to exist.

A natural extension of the classic Menelaus theorem is to consider other configurations (combinations of lines and their intersection points). Generalized Menelaus theorem (a line crossing a polygon) that corresponds to the wheel graph is well known. Below we prove the smallest possible "bipartite" Menelaus theorem.

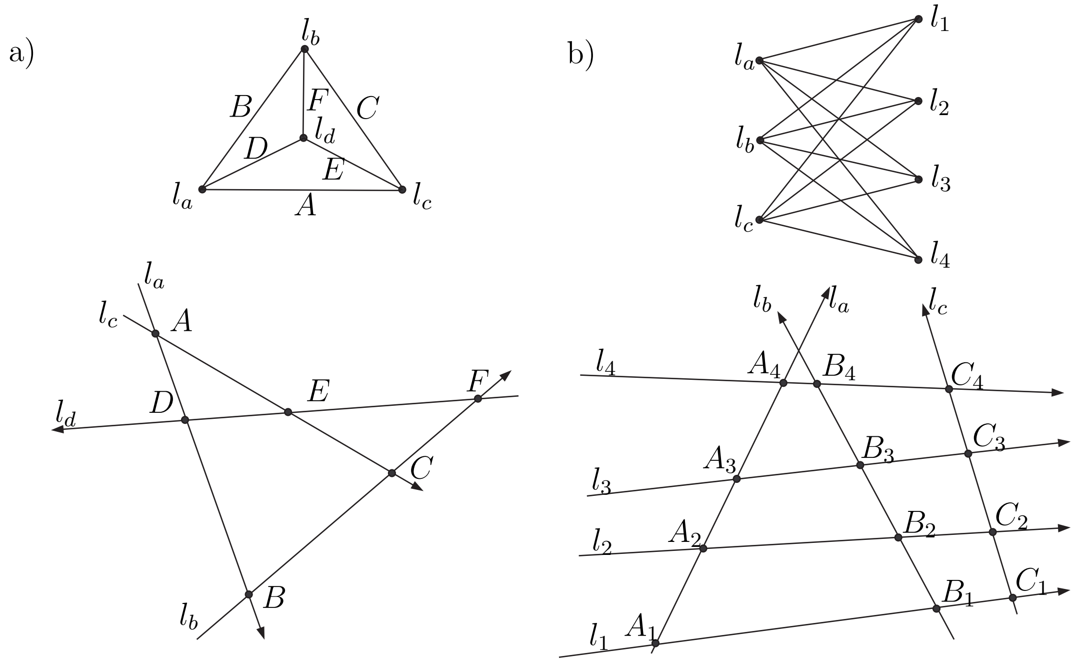

Figure 4: Classic and bipartite Menelaus' theorems and their intersection graphs.

Theorem 11 (Bipartite Menelaus) *Consider seven directed lines $l_{a,b,c}, l_{1,2,3,4}$ forming an irregular grid (see Fig. 4b) and their intersection points A_i, B_i, C_i. Then the following holds:*

$$\begin{aligned} & A_1A_3 \cdot A_2C_2 \cdot A_4C_4 \cdot B_1C_1 \cdot B_2B_4 \cdot B_3C_3 + A_2A_4 \cdot A_1C_1 \cdot A_3C_3 \cdot B_1B_3 \cdot B_2C_2 \cdot B_4C_4 \\ = & A_3A_4 \cdot A_1C_1 \cdot A_2C_2 \cdot B_1B_2 \cdot B_3C_3 \cdot B_4C_4 + A_2A_3 \cdot A_1C_1 \cdot A_4C_4 \cdot B_1B_4 \cdot B_2C_2 \cdot B_3C_3 \\ + & A_1A_2 \cdot A_3C_3 \cdot A_4C_4 \cdot B_1C_1 \cdot B_2C_2 \cdot B_3B_4 + A_1A_4 \cdot A_2C_2 \cdot A_3C_3 \cdot B_1C_1 \cdot B_2B_3 \cdot B_4C_4. \end{aligned} \tag{6}$$

The segments are signed as before.

Proof Add to the drawing the x axis, which is neither parallel nor orthogonal to any of the seven lines. For $i \in \{1, 2, 3, 4\}$ let x_i, y_i, z_i be the projections of A_i, B_i, C_i respectively onto this axis. Then the theorem's claim (6) can be rewritten as follows:

$$\begin{aligned} & (x_1-x_3)(x_2-z_2)(x_4-z_4)(y_1-z_1)(y_2-y_4)(y_3-z_3) + (x_2-x_4)(x_1-z_1)(x_3-z_3)(y_1-y_3)(y_2-z_2)(y_4-z_4) \\ = & (x_3-x_4)(x_1-z_1)(x_2-z_2)(y_1-y_2)(y_3-z_3)(y_4-z_4) + (x_2-x_3)(x_1-z_1)(x_4-z_4)(y_1-y_4)(y_2-z_2)(y_3-z_3) \\ + & (x_1-x_2)(x_3-z_3)(x_4-z_4)(y_1-z_1)(y_2-z_2)(y_3-y_4) + (x_1-x_4)(x_2-z_2)(x_3-z_3)(y_1-z_1)(y_2-y_3)(y_4-z_4). \end{aligned}$$

With a help of a symbolic calculator it is easy to verify that the above formula is equivalent to:

$$\det \begin{pmatrix} x_1 - y_1 & x_1 - z_1 & z_1(x_1 - y_1) & y_1(x_1 - z_1) \\ x_2 - y_2 & x_2 - z_2 & z_2(x_2 - y_2) & y_2(x_2 - z_2) \\ x_3 - y_3 & x_3 - z_3 & z_3(x_3 - y_3) & y_3(x_3 - z_3) \\ x_4 - y_4 & x_4 - z_4 & z_4(x_4 - y_4) & y_4(x_4 - z_4) \end{pmatrix} = 0. \tag{7}$$

Observe that (7) is invariant under all configuration-preserving permutations of the lines $l_{a,b,c}$ and $l_{1,2,3,4}$, which is less than obvious given only the original statement.

The proof consists of two substantially different cases.

Case I. There exist two lines, say, l_a and l_b, so that $A_i \neq B_i$ for all $i \in \{1, 2, 3, 4\}$. By appropriately scaling and translating the y axis we can make $l_a = \{x = y\}$. The rest of the proof will be done in the homogenous coordinates. Let $l_a = (1 : -1 : 0)$ and $l_b = (\alpha : \beta : \gamma)$. Since l_a and l_b are not equal, either $\gamma \neq 0$ or $\alpha \neq -\beta$. For $i \in \{1, 2, 3, 4\}$ the intersection point of l_a and l_i projects to x_i, and therefore $A_i = l_a \cap (1 : 0 : -x_i) = (x_i : x_i : 1)$. Likewise, $B_i = (\beta y_i : -\gamma - \alpha y_i : \beta)$. The two (distinct!) points uniquely define l_i:

$$l_i = (x_i\beta + \gamma + \alpha y_i : -x_i\beta + \beta y_i : -\gamma x_i - \alpha x_i y_i - \beta x_i y_i).$$

The four lines l_i for $i \in \{1, 2, 3, 4\}$ intersect with the vertical lines $(1 : 0 : -z_i)$ at the following points that must be collinear (as they all lie on l_c):

$$C_i = (\beta z_i(x_i - y_i) : \gamma(z_i - x_i) - \beta x_i(y_i - z_i) + \alpha y_i(z_i - x_i) : \beta(x_i - y_i)). \tag{8}$$

The points are collinear if and only if the 3×4 matrix whose ith line is (8) has rank less than 3. Rank-preserving transformation of these matrix reduce it to the matrix M with the following ith line (we may divide by $\beta \neq 0$ because l_b is not vertical):

$$z_i(x_i - y_i), (z_i - x_i)(\gamma + (\alpha + \beta)y_i), x_i - y_i.$$

Since matrix M has rank less than 3 for some α, β, and γ subject to the condition that either $\gamma \neq 0$ or $\alpha \neq -\beta$, it is equivalent to the following 4×4 matrix being singular, which implies (7):

$$\begin{pmatrix} z_1(x_1 - y_1) & z_1 - x_1 & (z_1 - x_1)y_1 & x_1 - y_1 \\ z_2(x_2 - y_2) & z_2 - x_2 & (z_2 - x_2)y_2 & x_2 - y_2 \\ z_3(x_3 - y_3) & z_3 - x_3 & (z_3 - x_3)y_3 & x_3 - y_3 \\ z_4(x_4 - y_4) & z_4 - x_4 & (z_4 - x_4)y_4 & x_4 - y_4 \end{pmatrix}.$$

Case II. For any two lines from l_a, l_b, l_c there is some l_i going through their intersection. Reorder the lines so that $A_1 = C_1$, $B_2 = C_2$, and $A_3 = B_3$ (see Figure 5).

Plugging the identities into (6) we reduce it to

$$A_1A_3 \cdot B_2B_4 \cdot A_4C_4 = B_2B_3 \cdot B_4C_4 \cdot A_1A_4,$$

which is true by classic Menelaus' theorem. □

We note that Theorem 11 is the "minimal" projective theorem that holds for the intersection points of two sets of lines in general position. Indeed, it follows from the proof that

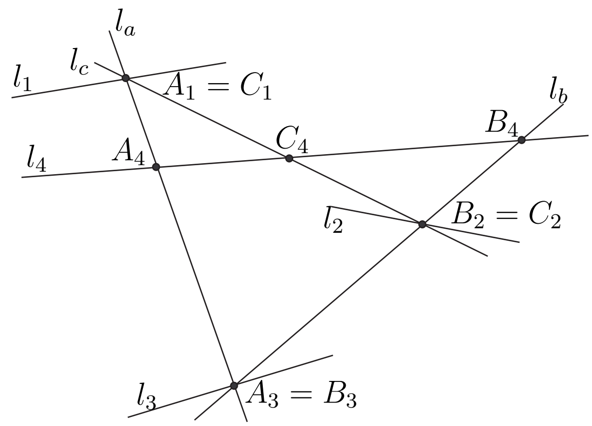

Figure 5: Degenerate case.

for any assignment of the eleven variables $x_{1,2,3,4}$, $y_{1,2,3,4}$, $z_{1,2,3}$ there is a configuration of lines whose intersection points project to those variables. This is because there always exist α, β, γ such that $\gamma \neq 0$ or $\alpha + \beta \neq 0$ and C_1, C_2, C_3 as defined in (8) are collinear. Hence no projection-invariant equation can be imposed on the lengths of the segments that would not include all the twelve points. Other configurations with as many or fewer intersection points are of no use either: six lines intersecting two can project to any collection of twelve points.

To highlight the projective nature of Theorem 11, we may rewrite (6) in a form that is invariant under projection:

$$\begin{aligned} 0 = {} & \frac{A_3A_4}{A_1A_4} \cdot \frac{B_1B_2}{B_1B_4} \cdot \frac{B_3C_3}{A_3C_3} \cdot \frac{B_4C_4}{A_4C_4} - \frac{A_2A_4}{A_1A_4} \cdot \frac{B_1B_3}{B_1B_4} \cdot \frac{B_2C_2}{A_2C_2} \cdot \frac{B_4C_4}{A_4C_4} + \frac{A_2A_3}{A_1A_4} \cdot \frac{B_2C_2}{A_2C_2} \cdot \frac{B_3C_3}{A_3C_3} \\ & + \frac{A_1A_2}{A_1A_4} \cdot \frac{B_3B_4}{B_1B_4} \cdot \frac{B_1C_1}{A_1C_1} \cdot \frac{B_2C_2}{A_2C_2} - \frac{A_1A_3}{A_1A_4} \cdot \frac{B_2B_4}{B_1B_4} \cdot \frac{B_1C_1}{A_1C_1} \cdot \frac{B_3C_3}{A_3C_3} + \frac{B_2B_3}{B_1B_4} \cdot \frac{B_1C_1}{A_1C_1} \cdot \frac{B_4C_4}{A_4C_4}, \end{aligned}$$

where $\frac{AB}{CD}$ for parallel AB and CD equals $\frac{|AB|}{|CD|}$ when the two segments have the same direction, and $-\frac{|AB|}{|CD|}$ otherwise.

Finally, we transform (7) to draw an analogy with (4):

$$\det \begin{pmatrix} A_1B_1 & A_1C_1 & C_1C_1 \cdot A_1B_1 & B_1B_1 \cdot A_1C_1 \\ A_2B_2 & A_2C_2 & C_2C_1 \cdot A_2B_2 & B_2B_1 \cdot A_2C_2 \\ A_3B_3 & A_3C_3 & C_3C_1 \cdot A_3B_3 & B_3B_1 \cdot A_3C_3 \\ A_4B_4 & A_4C_4 & C_4C_1 \cdot A_4B_4 & B_4B_1 \cdot A_4C_4 \end{pmatrix} = 0. \tag{9}$$