The Discrete Logarithm Problem with Preprocessing

<u>Henry Corrigan–Gibbs</u> and Dmitry Kogan Stanford University

> Eurocrypt – 1 May 2018 Tel Aviv, Israel











The discrete-log problem



The discrete-log problem



Why do we believe this problem is hard?

Generic lower bounds give us confidence

Theorem. [Shoup'97] Every **generic** discrete-log algorithm that

- operates in a group of prime order N and
- succeeds with probability at least $\frac{1}{2}$

must run in time $\Omega(N^{1/2})$.

Generic lower bounds give us confidence

Theorem. [Shoup'97] Every generic discrete-log algorithm that

- operates in a group of prime order N and
- succeeds with probability at least $\frac{1}{2}$

must run in time $\Omega(N^{1/2})$.

Generic attack in 256–bit group takes $\approx 2^{128}$ time.

Generic lower bounds give us confidence

Theorem. [Shoup'97] Every generic discrete-log algorithm that
operates in a group of prime order N and

• succeeds with probability at least $\frac{1}{2}$

must run in time $\Omega(N^{1/2})$.

Generic attack in 256-bit group takes $\approx 2^{128}$ time.

Best attacks on standard EC groups are generic Generic algorithms can only make "black-box" use of the group operation

Generic-group model:

- Group is defined by an injective "labeling" function $\sigma: \mathbb{Z}_N \to \{0,1\}^*$
- Algorithm has access to a group-operation oracle: $\mathcal{O}_{\sigma}(\sigma(i), \sigma(j)) \mapsto \sigma(i+j)$

[Nechaev'94], [Shoup'97], [Maurer'05]

Generic dlog algorithm takes as input $(\sigma(1), \sigma(x))$, representing (g, g^x) , make queries to \mathcal{O}_{σ} , outputs x. [Measure running time by query complexity] Generic algorithms can only make "black-box" use of the group operation

Very useful way to

understand hardness

[BB04,B05,M05,D06,

B08,Y15,...]

49

Generic-group model:

- Group is defined by an injective "labeling" function $\sigma: \mathbb{Z}_N \to \{0,1\}^*$
- Algorithm has access to a group-operation oracle: $\mathcal{O}_{\sigma}(\sigma(i), \sigma(j)) \mapsto \sigma(i+j)$

Generic dlog algorithm takes as input ($\sigma(1 (g, g^x))$, make queries to \mathcal{O}_{σ} , outputs x. [Measure running time by query complet

- Premise of generic-group model: the adversary knows nothing about the structure of the group G in advance
- In reality: the adversary knows a lot about G!
 - ≻ G is one of a small number of groups: NIST P-256, Curve25519, ...
- A realistic adversary can perform G-specific preprocessing!
- Existing generic-group lower bounds say <u>nothing</u> about preprocessing attacks! [H80, Yao90, FN91, ...]

- Premise of generic-group model: the adversary knows nothing about the structure of the group G in advance
- In reality: the adversary knows a lot about **G**!
 - ➤ G is one of a small number of groups: NIST P-256, Curve25519, ...
- A realistic adversary can perform G-specific preprocessing!
- Existing generic-group lower bounds say <u>nothing</u> about preprocessing attacks! [H80, Yao90, FN91, ...]

- Premise of generic-group model: the adversary knows nothing about the structure of the group G in advance
- In reality: the adversary knows a lot about **G**!
 - ➤ G is one of a small number of groups: NIST P-256, Curve25519, ...
- A realistic adversary can perform G-specific preprocessing!
- Existing generic-group lower bounds say <u>nothing</u> about preprocessing attacks! [H80, Yao90, FN91, ...]

- Premise of generic-group model: the adversary knows nothing about the structure of the group G in advance
- In reality: the adversary knows a lot about **G**!
 - ➤ G is one of a small number of groups: NIST P-256, Curve25519, ...
- A realistic adversary can perform G-specific preprocessing!
- Existing generic-group lower bounds say <u>nothing</u> about preprocessing attacks! [H80, Yao90, FN91, ...]















Rest of this talk

Background: Preprocessing attacks are relevant

• Preexisting $S = T = \tilde{O}(N^{1/3})$ generic attack on discrete log

Our results: Preprocessing lower-bounds and attacks

- The $\tilde{O}(N^{1/3})$ generic dlog attack is optimal
- Any such attack must use <u>lots</u> of preprocessing: $\Omega(N^{2/3})$
- New $\tilde{O}(N^{1/5})$ preprocessing attack on DDH-like problem

Open questions

Rest of this talk

Background: Preprocessing attacks are relevant

• Preexisting $S = T = \tilde{O}(N^{1/3})$ generic attack on discrete log

Our results: Preprocessing lower-bounds and attacks • The $\tilde{O}(N^{1/3})$ generic dlog attack is optimal

- Any such attack must use <u>lots</u> of preprocessing: $\Omega(N^{2/3})$
- New $\tilde{O}(N^{1/5})$ preprocessing attack on DDH-like problem

Open questions

A preexisting result...

Theorem. [Mihalcik 2010] [Lee, Cheon, Hong 2011] [Bernstein and Lange 2013]

There is a generic dlog algorithm with preprocessing that:

- uses S bits of group-specific advice,
- uses T online time, and
- succeeds with probability ϵ , such that:

$$ST^2 = \tilde{O}(\epsilon N)$$

.... building on prior work on multiple-discrete-log algorithms [ESST99,KS01,HMCD04,BL12] A preexisting result...

Theorem. [Mihalcik 2010] [Lee, Cheon, Hong 2011] [Bernstein and Lange 2013] There is a generic dlog algorithm with preprocessing that:

 $ST^2 = \tilde{O}(\epsilon N)$

- uses S bits of group-specific advice,
- uses T online time, and
- succeeds with probability ϵ , such that:

Will sketch the algorithm for $S = T = N^{1/3}$, constant ϵ .

.... building on prior work on multiple-discrete-log algorithms [ESST99,KS01,HMCD04,BL12] A preexisting result...

Theorem. [Mihalcik 2010] [Lee, Cheon, Hong 2011] [Bernstein and Lange 2013]

There is a generic dlog algorithm with preprocessing that:

- uses S bits of group-specific advice,
- uses T online time, and
- succeeds with probability ϵ , such that:

$$ST^2 = \tilde{O}(\epsilon N)$$

.... building on prior work on multiple-discrete-log algorithms [ESST99,KS01,HMCD04,BL12]

Define a pseudo-random walk on G:

$$g^x \mapsto g^{x+\alpha}$$
 where $\alpha = \text{Hash}(g^x)$ is a random function

Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



Define a pseudo-random walk on G:

$$g^x \mapsto g^{x+\alpha}$$
 where $\alpha = \text{Hash}(g^x)$ is a random function

$$g^x \qquad g^{x+\alpha_1}$$

Define a pseudo-random walk on G:

$$g^x \mapsto g^{x+\alpha}$$
 where $\alpha = \text{Hash}(g^x)$ is a random function



Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



If you know the dlog of the endpoint of a walk, you know the dlog of the starting point!
Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



Define a pseudo-random walk on G:

 $g^x \mapsto g^{x+\alpha}$ where $\alpha = \text{Hash}(g^x)$ is a random function



- Build $N^{1/3}$ chains of length $N^{1/3}$
- Store dlogs of chain endpoints

- Walk $O(N^{1/3})$ steps
- When you hit a stored point, output the discrete log



- Build $N^{1/3}$ chains of length $N^{1/3}$
- Store dlogs of chain endpoints

Advice: $\tilde{O}(N^{1/3})$ bits

- Walk $O(N^{1/3})$ steps
- When you hit a stored point, output the discrete log



- Build $N^{1/3}$ chains of length $N^{1/3}$
- Store dlogs of chain endpoints

Advice: $\tilde{O}(N^{1/3})$ bits

- Walk $O(N^{1/3})$ steps
- When you hit a stored point, output the discrete log



- Build $N^{1/3}$ chains of length $N^{1/3}$
- Store dlogs of chain endpoints

Advice: $\tilde{O}(N^{1/3})$ bits

- Walk $O(N^{1/3})$ steps
- When you hit a stored point, output the discrete log



- Build $N^{1/3}$ chains of length $N^{1/3}$
- Store dlogs of chain endpoints

Advice: $\tilde{O}(N^{1/3})$ bits

- Walk $O(N^{1/3})$ steps
- When you hit a stored point, output the discrete log



- Build $N^{1/3}$ chains of length $N^{1/3}$
- Store dlogs of chain endpoints

Advice: $\tilde{O}(N^{1/3})$ bits

- Walk $O(N^{1/3})$ steps
- When you hit a stored point, output the discrete log



- Build $N^{1/3}$ chains of length $N^{1/3}$
- Store dlogs of chain endpoints

Advice: $\tilde{O}(N^{1/3})$ bits

- Walk $O(N^{1/3})$ steps
- When you hit a stored point, output the discrete log



- Build $N^{1/3}$ chains of length $N^{1/3}$
- Store dlogs of chain endpoints

Advice: $\tilde{O}(N^{1/3})$ bits

Online phase

- Walk $O(N^{1/3})$ steps
- When you hit a stored point, output the discrete log

Time: $\tilde{O}(N^{1/3})$ steps



[M10, LCH11, BL13]

- Build $N^{1/3}$ chains of length $N^{1/3}$
- Store dlogs of chain endpoints

Advice: $\tilde{O}(N^{1/3})$ bits

Online phase

- Walk $O(N^{1/3})$ steps
- When you hit a stored point, output the discrete log

Time: $\tilde{O}(N^{1/3})$ steps



[M10, LCH11, BL13]

Generic discrete log

- \rightarrow Without preprocessing: $\Omega(N^{1/2})$
- \rightarrow With preprocessing: $\tilde{O}(N^{1/3})$

256-bit ECDL 2¹²⁸ time 2⁸⁶ time

Related preprocessing attacks break:

- Multiple discrete log problem
- One-round Even-Mansour cipher
- Merkle–Damgård hash with random IV

[This paper] [FJM14] [CDGS17]

Generic discrete log

- \rightarrow Without preprocessing: $\Omega(N^{1/2})$
- \rightarrow With preprocessing: $\tilde{O}(N^{1/3})$

<u>256-bit ECDL</u> 2¹²⁸ time

2⁸⁶ time

Related preprocessing attacks break:

- Multiple discrete log problem
- One-round Even-Mansour cipher
- Merkle–Damgård hash with random IV

[This paper] [FJM14] [CDGS17]



Generic discrete log

- \rightarrow Without preprocessing: $\Omega(N^{1/2})$
- \rightarrow With preprocessing: $\tilde{O}(N^{1/3})$

<u>256-bit ECDL</u> 2¹²⁸ time

2⁸⁶ time

Related preprocessing attacks break:

- Multiple discrete log problem
- One-round Even-Mansour cipher
- Merkle–Damgård hash with random IV

[This paper] [FJM14] [CDGS17]









This talk

Background: Preprocessing attacks are relevant

• Preexisting $S = T = \tilde{O}(N^{1/3})$ generic attack on discrete log

Our results: Preprocessing lower-bounds and attacks

- The $\tilde{O}(N^{1/3})$ generic dlog attack is optimal
- Any such attack must use <u>lots</u> of preprocessing: $\Omega(N^{2/3})$
- New $\tilde{O}(N^{1/5})$ preprocessing attack on DDH-like problem

Open questions

Every generic dlog algorithm with preprocessing that:

- uses *S* bits of group-specific advice,
- uses T online time, and
- succeeds with probability ϵ , must satisfy:

$$ST^2 = \widetilde{\Omega}(\epsilon N)$$

Every generic dlog algorithm with preprocessing that:

- uses *S* bits of group-specific advice,
- uses T online time, and
- succeeds with probability ϵ , must satisfy:

$$ST^2 = \widetilde{\Omega}(\epsilon N)$$

This bound is tight for the full range of parameters (up to log factors)

Every generic dlog algorithm with preprocessing that:

- uses *S* bits of group-specific advice,
- uses T online time, and
- succeeds with probability ϵ , must satisfy:

$$ST^2 = \widetilde{\Omega}(\epsilon N)$$

Every generic dlog algorithm with preprocessing that:

- uses *S* bits of group-specific advice,
- uses T online time, and
- succeeds with probability ε, must satisfy:

$$ST^2 = \widetilde{\Omega}(\epsilon N)$$

Shoup's proof technique (1997) relies on ${\cal A}$ having no information about the group ${\mathbb G}$ when it starts running

→ Need different proof technique

Every generic dlog algorithm with preprocessing that:

- uses *S* bits of group-specific advice,
- uses T online time, and
- succeeds with probability ϵ , must satisfy:

$$ST^2 = \widetilde{\Omega}(\epsilon N)$$

Every generic dlog algorithm with preprocessing that:

- uses *S* bits of group-specific advice,
- uses T online time, and
- succeeds with probability ε, must satisfy:

$$ST^2 = \widetilde{\Omega}(\epsilon N)$$

Theorem. [Our paper] Furthermore, the preprocessing time P must satisfy $PT + T^2 = \Omega(\epsilon N)$

Every generic dlog algorithm with preprocessing that:

- uses *S* bits of group-specific advice,
- uses T online time, and
- succeeds with probability ϵ , must satisfy:

$$ST^2 = \widetilde{\Omega}(\epsilon N)$$

Online time $N^{1/3}$ implies $\Omega(N^{2/3})$ preprocessing

Theorem. [Our paper]

Furthermore, the preprocessing time *P* must satisfy $PT + T^2 = \Omega(\epsilon N)$

Every generic dlog algorithm with preprocessing that:

- uses *S* bits of group-specific advice,
- uses T online time, and
- succeeds with probability ε, must satisfy:

$$ST^2 = \widetilde{\Omega}(\epsilon N)$$

Theorem. [Our paper] Furthermore, the preprocessing time P must satisfy $PT + T^2 = \Omega(\epsilon N)$

Reminder: Generic-group model

- A group is defined by an injective "labeling" function $\sigma: \mathbb{Z}_N \to \{0,1\}^*$
- Algorithm has access to a group-operation oracle: $\mathcal{O}_{\sigma}(\sigma(i), \sigma(j)) \mapsto \sigma(i+j)$

E.g., A dlog algorithm takes as input $(\sigma(1), \sigma(x))$, representing (g, g^x) , make queries to \mathcal{O}_{σ} , outputs x.

We prove the lower bound using an incompressibility argument [Yao90, GT00, DTT10, DHT12, DGK17...]

Use \mathcal{A} to compress the mapping $\sigma: \mathbb{Z}_N \to \{0,1\}^*$ that defines the group Similar technique used in [DHT12]



• Adv \mathcal{A} uses advice S and online time T such that $ST^2 = o(N)$ \Rightarrow Encoder compresses well

• Random string is incompressible \Rightarrow Lower bound on *S* and *T*

We prove the lower bound using an incompressibility argument [Yao90, GT00, DTT10, DHT12, DGK17...]

Use \mathcal{A} to compress the mapping $\sigma: \mathbb{Z}_N \to \{0,1\}^*$ that defines the group Similar technique used in [DHT12]



- Adv \mathcal{A} uses advice S and online time T such that $ST^2 = o(N)$ \Rightarrow Encoder compresses well
- Random string is incompressible \Rightarrow Lower bound on *S* and *T*
We prove the lower bound using an incompressibility argument [Yao90, GT00, DTT10, DHT12, DGK17...]

Use \mathcal{A} to compress the mapping $\sigma: \mathbb{Z}_N \to \{0,1\}^*$ that defines the group Similar technique used in [DHT12]



• Adv \mathcal{A} uses advice S and online time T such that $ST^2 = o(N)$

 \Rightarrow Encoder compresses well

• Random string is incompressible \Rightarrow Lower bound on S and T

We prove the lower bound using an incompressibility argument [Yao90, GT00, DTT10, DHT12, DGK17...]

Use \mathcal{A} to compress the mapping $\sigma: \mathbb{Z}_N \to \{0,1\}^*$ that defines the group





We prove the lower bound using an incompressibility argument [Yao90, GT00, DTT10, DHT12, DGK17...]

Use \mathcal{A} to compress the mapping $\sigma: \mathbb{Z}_N \to \{0,1\}^*$ that defines the group Similar technique used in [DHT12]



• Adv \mathcal{A} uses advice S and online time T such that $ST^2 = o(N)$

 \Rightarrow Encoder compresses well

• Random string is incompressible \Rightarrow Lower bound on S and T

Encoder



































- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"


Proof idea: Use preprocessing dlog adversary $(\mathcal{A}_0, \mathcal{A}_1)$ to build a compressed representation of the mapping σ . [Yao90, GT00, DHT12]

- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



Proof idea: Use preprocessing dlog adversary $(\mathcal{A}_0, \mathcal{A}_1)$ to build a compressed representation of the mapping σ . [Yao90, GT00, DHT12]

- Run \mathcal{A}_1 on I instances
- Whenever \mathcal{A}_1 outputs a dlog, we get one value $\sigma(i)$ "for free"



Easy case: The response to all of A_1 's queries are distinct • A_1 outputs a discrete log "for free" \Rightarrow Compress by $\approx \log N$ bits

Harder case:The response to query t is the same as the
response to query t' < t.

- A naïve encoding "pays twice" for the same value $\sigma(i) \Rightarrow$ No savings \otimes
- Instead, encoder writes a pointer to query t'If the encoder runs \mathcal{A}_1 on I instances, requires $\log IT + \log T$ bits.

Pointer toIndex ofquery t'query t

Easy case: The response to all of \mathcal{A}_1 's queries are distinct

• \mathcal{A}_1 outputs a discrete log "for free" \Rightarrow Compress by $\approx \log N$ bits

Harder case:The response to query t is the same as the
response to query t' < t.

- A naïve encoding "pays twice" for the same value $\sigma(i) \Rightarrow$ No savings \otimes
- Instead, encoder writes a pointer to query t'If the encoder runs \mathcal{A}_1 on I instances, requires $\log IT + \log T$ bits.

Pointer toIndex ofquery t'query t

Easy case: The response to all of \mathcal{A}_1 's queries are distinct

• \mathcal{A}_1 outputs a discrete log "for free" \Rightarrow Compress by $\approx \log N$ bits

Harder case: The response to query t is the same as the response to query t' < t.

- A naïve encoding "pays twice" for the same value $\sigma(i) \Rightarrow$ No savings \otimes
- Instead, encoder writes a pointer to query t'If the encoder runs \mathcal{A}_1 on I instances, requires $\log IT + \log T$ bits.

Pointer to Index of

query t' query t

Easy case: The response to all of \mathcal{A}_1 's queries are distinct

- \mathcal{A}_1 outputs a discrete log "for free" \Rightarrow Compress by $\approx \log N$ bits
- **Harder case:** The response to query t is the same as the response to query t' < t.
- A naïve encoding "pays twice" for the same value $\sigma(i) \Rightarrow$ No savings \otimes

• Instead, encoder writes a pointer to query t'If the encoder runs \mathcal{A}_1 on I instances, requires $\log IT + \log T$ bits. Pointer to Index of

query t' query t

Easy case: The response to all of \mathcal{A}_1 's queries are distinct

• \mathcal{A}_1 outputs a discrete log "for free" \Rightarrow Compress by $\approx \log N$ bits

Harder case: The response to query t is the same as the response to query t' < t.

- A naïve encoding "pays twice" for the same value $\sigma(i) \Rightarrow$ No savings \otimes
- Instead, encoder writes a pointer to query t'If the encoder runs \mathcal{A}_1 on I instances, requires $\log IT + \log T$ bits.

Pointer to
query t'Index of
query t

Easy case: The response to all of \mathcal{A}_1 's queries are distinct

• \mathcal{A}_1 outputs a discrete log "for free" \Rightarrow Compress by $\approx \log N$ bits

Harder case: The response to query t is the same as the response to query t' < t.

- A naïve encoding "pays twice" for the same value $\sigma(i) \Rightarrow$ No savings \otimes
- Instead, encoder writes a pointer to query t'If the encoder runs \mathcal{A}_1 on I instances, requires $\log IT + \log T$ bits.

Pointer to query *t*'

Index of query *t* Each execution of A_1 saves at least 1 bit, when: $\log IT^2 < \log N$, or $I < N/T^2$

Completing the proof

- We run the adversary \mathcal{A}_1 on $I = N/T^2$ instances
- Each execution compresses by \geq 1 bit
- BUT, we have to include the *S*-bit advice string in the encoding

Encoding overhead =
$$S - \frac{N}{T^2} \ge 0 \implies ST^2 = \Omega(N)$$

Extra complications

- Algorithms that succeed on an ϵ -fraction of group elements
 - Use the random self-reducibility of dlog
 - Hardcode a good set of random coins for \mathcal{A}_1 into $\mathsf{Enc}(\sigma)$

- Decisional type problems (DDH, etc.)
 - $-A_1$ only outputs 1 bit—prior argument fails because encoding the runtime in $\log T$ bits is too expensive
 - Run \mathcal{A}_1 on batches of inputs

[See paper for details]

| | Upper bound | Lower bound | Time T | |
|---------------|--------------------------------|---|----------------|-----------------------|
| Discrete log: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | - For |
| CDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | $\epsilon = N^{-1/4}$ |
| DDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon^2 N)$ | $\leq N^{1/4}$ | $S = N^{1/4}$ |
| | | | $> N^{1/8}$ | |

| | Upper bound | Lower bound | Time T | |
|---------------|--------------------------------|---|----------------|-----------------------|
| Discrete log: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | - For |
| CDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | $\epsilon = N^{-1/4}$ |
| DDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon^2 N)$ | $\leq N^{1/4}$ | $S = N^{1/4}$ |
| | | | $> N^{1/8}$ | |

| | Upper bound | Lower bound | Time T | |
|---------------|--------------------------------|---|----------------|-----------------------|
| Discrete log: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | - For |
| CDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | $\epsilon = N^{-1/4}$ |
| DDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon^2 N)$ | $\leq N^{1/4}$ | $S = N^{1/4}$ |
| | | | $\geq N^{1/8}$ | |

| | Upper bound | Lower bound | Time T | |
|---------------|--------------------------------|---|----------------|-----------------------|
| Discrete log: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | - For |
| CDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | $\epsilon = N^{-1/4}$ |
| DDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon^2 N)$ | $\leq N^{1/4}$ | $S = N^{1/4}$ |
| | | | $\geq N^{1/8}$ | |
| | | | | |
| | | | Better | attack? |

| | Upper bound | Lower bound | Time T | |
|---------------|--------------------------------|---|----------------|-----------------------|
| Discrete log: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | - For |
| CDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | $\epsilon = N^{-1/4}$ |
| DDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon^2 N)$ | $\leq N^{1/4}$ | $S = N^{1/4}$ |
| | | | $\geq N^{1/8}$ | |

| | Upper bound | Lower bound | Time T | |
|---------------|----------------------------------|---|----------------|-----------------------|
| Discrete log: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | - For |
| CDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | $\epsilon = N^{-1/4}$ |
| DDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon^2 N)$ | $\leq N^{1/4}$ | $S = N^{1/4}$ |
| | | | $\geq N^{1/8}$ | |
| sqDDH: | $ST^2 = \tilde{O}(\epsilon^2 N)$ | $ST^2 = \widetilde{\Omega}(\epsilon^2 N)$ | $N^{1/8}$ | |

| | Upper bound | Lower bound | Time T | |
|---------------|----------------------------------|---|----------------|-----------------------|
| Discrete log: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | For |
| CDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon N)$ | $N^{1/4}$ | $\epsilon = N^{-1/4}$ |
| DDH: | $ST^2 = \tilde{O}(\epsilon N)$ | $ST^2 = \widetilde{\Omega}(\epsilon^2 N)$ | $\leq N^{1/4}$ | $S = N^{1/4}$ |
| | | | $\geq N^{1/8}$ | |
| sqDDH: | $ST^2 = \tilde{O}(\epsilon^2 N)$ | Our new results | $N^{1/8}$ | |

Definition. The sqDDH problem is to distinguish $(g, g^x, g^{(x^2)})$ from (g, g^x, g^y)

Why it's interesting:

• For generic online-only algs,

• For generic preprocessing algs,

it's as hard as discrete log we show that it's "much easier"

for $x, y \leftarrow_R \mathbb{Z}_N$.

Definition. The sqDDH problem is to distinguish $(g, g^x, g^{(x^2)})$ from (g, g^x, g^y) for $x, y \leftarrow_R \mathbb{Z}_N$.

Why it's interesting:

• For generic online-only algs,

it's as hard as discrete log

• For generic preprocesssing algs,

we show that it's "much easier"

Definition. The sqDDH problem is to distinguish $(g, g^x, g^{(x^2)})$ from (g, g^x, g^y) for $x, y \leftarrow_R \mathbb{Z}_N$.

Why it's interesting:

- For generic online-only algs,
- For generic preprocessing algs,

it's as hard as discrete log we show that it's "much easier"

Definition. The sqDDH problem is to distinguish $(g, g^x, g^{(x^2)})$ from (g, g^x, g^y) for $x, y \leftarrow_R \mathbb{Z}_N$.

Why it's interesting:

- For generic online-only algs,
- For generic preprocesssing algs,

it's as hard as discrete log we show that it's "much easier"

This talk

Background: Preprocessing attacks are relevant

• Preexisting $S = T = \tilde{O}(N^{1/3})$ generic attack on discrete log

Our results: Preprocessing lower-bounds and attacks • The $\tilde{O}(M^{1/3})$ generic dlog attack is optimal

- The $\tilde{O}(N^{1/3})$ generic dlog attack is optimal
- Any such attack must use <u>lots</u> of preprocessing: $\Omega(N^{2/3})$
- New $\tilde{O}(N^{1/5})$ preprocessing attack on DDH-like problem

Open questions

Open questions and recent progress

- Tightness of DDH upper/lower bounds?
 - Is it as hard as dlog or as easy as sqDDH?
- Non-generic preprocessing attacks on ECDL?
 - As we have for \mathbb{Z}_p^*

Coretti, Dodis, and Guo (2018)

- Elegant proofs of generic-group lower bounds using "presampling" (à la Unruh, 2007)
- Prove hardness of "one-more" dlog, KEA assumptions, ...

This talk

Background: Preprocessing attacks are relevant

• Preexisting $S = T = \tilde{O}(N^{1/3})$ generic attack on discrete log

Our results: Preprocessing lower-bounds and attacks

- The $\tilde{O}(N^{1/3})$ generic dlog attack is optimal
- Any such attack must use <u>lots</u> of preprocessing: $\Omega(N^{2/3})$
- New $\tilde{O}(N^{1/5})$ preprocessing attack on DDH-like problem

Open questions

This talk

Background: Preprocessing attacks are relevant

• Preexisting $S = T = \tilde{O}(N^{1/3})$ generic attack on discrete log

Our results: Preprocessing lower-bounds and attacks • The $\tilde{O}(N^{1/3})$ generic dlog attack is optimal

- Any such attack must use <u>lots</u> of preprocessing: $\Omega(N^{2/3})$
- New $\tilde{O}(N^{1/5})$ preprocessing attack on DDH-like problem

Open questions

Henry – henrycg@cs.stanford.edu Dima – dkogan@cs.stanford.edu https://eprint.iacr.org/2017/1113