DDH and ElGamal Encryption

Lecture 25: Key agreement from exponentiation

Lecture overview

Alice and Bob can compute g^{ab} after exchanging g^a and g^b. We ask when this shared group element looks random to an eavesdropper, then use it to encrypt a message.

Based on the instructor's Lecture 25, pp. 1–10 [Maj26].

Abstract and lecture roadmap

The **decisional Diffie–Hellman assumption** says that g^{ab} looks like a fresh random group element even when g^a and g^b are public.

1. Define the DDH challenge and fix our advantage convention.
2. Prove correctness and passive privacy of Diffie–Hellman key agreement.
3. Construct ElGamal encryption and prove privacy under DDH.
4. Compare DL, CDH, and DDH, and analyze the Legendre-symbol attack.

The question that drives the lecture

Does the difficulty of *computing* a secret imply that the secret *looks random*?

Source objectives: [Maj26, p. 2]; [BS23, §10.4–10.5]; [HPS14, §2.2–2.4].

The main conclusions

DDH gives passive key indistinguishability

The public transcript and real key (g^a, g^b, g^{ab}) are computationally indistinguishable from (g^a, g^b, U), where $U \xleftarrow{\$} \mathbb{G}$ is independent.

DDH gives ElGamal privacy

For messages in $\mathbb{G}$, the ciphertext $(g^b, m \cdot (g^a)^b)$ is IND-CPA secure under the DDH assumption. The proof replaces the shared mask by U.

The hardness implications go in one direction

$$\boxed{\text{DDH hardness} \implies \text{CDH hardness} \implies \text{DL hardness.}}$$

[BS23, Def. 10.8, p. 411; Ex. 11.5, p. 459; Lemmas 16.1–16.2, p. 687].

The group and the exponent space

Let $\mathbb{G} = \langle g \rangle$ be a finite cyclic group of known order q.

$$\mathbb{G} = \{1, g, g^2, \ldots, g^{q-1}\}, \qquad g^q = 1.$$

Group elements

$A, B, K, m \in \mathbb{G}$. We can multiply, invert, test equality, and exponentiate efficiently.

$$A \cdot B, \qquad A^{-1}, \qquad A^b.$$

Exponents

$a, b, r \in \mathbb{Z}_q$. Arithmetic on exponents is modulo q.

$$(g^a)^b = g^{ab}, \qquad g^a g^b = g^{a+b}.$$

Usually q is a large prime. The even-order group in the attack will be an explicit exception.

The source uses n for the group order. We use q and sample from $\{0, \ldots, q-1\}$, an equivalent complete residue system.

[Maj26, p. 3]; [HPS14, §2.2, pp. 64–67]; [BS23, §10.5, p. 408].

Two distributions with the same public pair

Sample $a, b, r \xleftarrow{\$} \mathbb{Z}_q$ independently.

Real Diffie–Hellman tuple

$$\mathcal{D}_0 = (g^a, g^b, g^{ab}).$$

The third element is determined by the first two group elements.

Random tuple

$$\mathcal{D}_1 = (g^a, g^b, g^r).$$

The third element is uniform on $\mathbb{G}$ and independent of the first two.

DDH intuition. An efficient observer cannot tell which distribution supplied the tuple. This is a computational claim: an unbounded observer can compute discrete logarithms and test the exponent relation.

[Maj26, pp. 3, 7]; [BS23, Attack Game 10.6, p. 410; Def. 10.8, p. 411].

The DDH guessing game

Public parameters: $(\mathbb{G}, q, g)$.

Challenger and adversary

1. Challenger samples $u \xleftarrow{\$} \{0,1\}$ and $a, b, r \xleftarrow{\$} \mathbb{Z}_q$.
2. It sends (A, B, T) to the adversary, where

$$A = g^a, \quad B = g^b, \quad T = \begin{cases} g^{ab} & u = 0, \\ g^r & u = 1. \end{cases}$$

3. The adversary returns a bit $\widehat{u}$ and wins if $\widehat{u} = u$.

The asymptotic assumption

For the chosen group family, every efficient adversary has negligible guessing advantage as the security parameter λ grows.

Source game: [Maj26, p. 7]; formal family-based formulation: [BS23, Def. 10.8, p. 411].

Two advantage conventions

Let D output 1 when it thinks the tuple is real. Put

$$p_0 = \Pr[D(\mathcal{D}_0) = 1], \qquad p_1 = \Pr[D(\mathcal{D}_1) = 1].$$

Distinguishing advantage used in this lecture

$$\boxed{\mathrm{Adv}_{\mathrm{DDH}}(D) = |p_0 - p_1|.}$$

This is the convention in Boneh–Shoup's DDH game.

If $p_0 \geq p_1$, guess $u = 0$ when $D = 1$ and $u = 1$ otherwise. Then

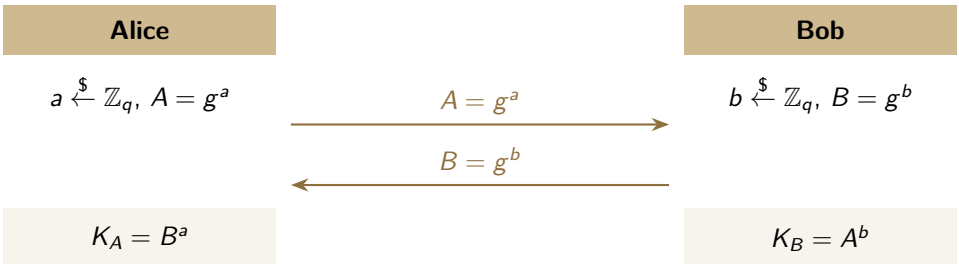
$$\Pr[\widehat{u} = u] = \tfrac{1}{2}p_0 + \tfrac{1}{2}(1 - p_1) = \tfrac{1}{2} + \tfrac{1}{2}\,\mathrm{Adv}_{\mathrm{DDH}}(D).$$

Guessing advantage $\left|\Pr[\widehat{u} = u] - \tfrac{1}{2}\right|$ is half the distinguishing advantage. The absolute-value identity also covers the opposite orientation.

[BS23, Attack Game 10.6, p. 410; Def. 10.8, p. 411]; compare the guessing game in source p. 7.

Diffie–Hellman key agreement

Public parameters: a cyclic group $\mathbb{G} = \langle g \rangle$ of order q.



Eavesdropper sees (g, A, B) and the public group description.

The two messages can be sent simultaneously: neither depends on the other.

[Maj26, p. 4]; [HPS14, Table 2.2, p. 68]; [BS23, §10.4.1, p. 407; Fig. 10.2, p. 408]; [DH76].

Correctness and privacy answer different questions

Correctness: do the honest parties agree?

$$\boxed{K_A = B^a = (g^b)^a = g^{ab} = (g^a)^b = A^b = K_B.}$$

This identity uses group arithmetic and holds for every $a, b \in \mathbb{Z}_q$.

Privacy proof idea. The transcript and test key together form exactly a DDH challenge.

Privacy against a passive observer, assuming DDH

$$(A, B, K) = (g^a, g^b, g^{ab}) \approx_c (g^a, g^b, U), \qquad U \xleftarrow{\$} \mathbb{G}.$$

Any efficient key-test distinguisher is already a DDH distinguisher with the same distinguishing advantage.

The shared object here is a group element. Deriving application keys as bit strings requires a suitable key-derivation step.

[HPS14, Table 2.2, p. 68]; [BS23, §10.4.1, p. 407; Def. 10.8, p. 411].

A small example in the quadratic residues modulo 23

Take $p = 23$, $q = 11$, and $\mathbb{G} = \langle 2 \rangle \subseteq \mathbb{F}_{23}^*$.

$$\mathbb{G} = \{1, 2, 4, 8, 16, 9, 18, 13, 3, 6, 12\}.$$

Alice chooses $a = 3$

$$A = 2^3 \equiv 8 \pmod{23}.$$

After receiving $B = 13$,

$$K_A = 13^3 \equiv 12 \pmod{23}.$$

Bob chooses $b = 7$

$$B = 2^7 \equiv 13 \pmod{23}.$$

After receiving $A = 8$,

$$K_B = 8^7 \equiv 12 \pmod{23}.$$

Both computations equal $2^{21} \equiv 2^{10} \equiv 12 \pmod{23}$ because exponents reduce modulo 11.

These small parameters illustrate the algebra. They provide no computational security.

Worked instance of [HPS14, Table 2.2, p. 68]; arithmetic checked independently.

The privacy theorem assumes a passive observer

Suppose Eve can replace messages. She intercepts $A = g^a$ and $B = g^b$, and sends g^e to Alice and g^f to Bob.

Alice's resulting key

$$K_A = (g^e)^a = A^e.$$

Eve knows e and the intercepted A, so she can compute K_A.

Bob's resulting key

$$K_B = (g^f)^b = B^f.$$

Eve knows f and the intercepted B, so she can compute K_B.

Eve can decrypt and relay messages between the two sessions. To resist this attack, the exchange must authenticate the public shares, for example using signatures and authenticated public keys.

[HPS14, Example 3.13, pp. 126–127; Fig. 3.1, p. 127].

From the exchange to public-key encryption

In this Diffie–Hellman exchange, Alice's first message can be published once and reused as a public key.

Receiver's persistent state

Choose $a \xleftarrow{\$} \mathbb{Z}_q$. Publish $A = g^a$ and keep a secret.

Sender's fresh contribution

For each message $m \in \mathbb{G}$, choose a fresh $b \xleftarrow{\$} \mathbb{Z}_q$ and compute

$$B = g^b, \qquad K = A^b, \qquad C = mK.$$

Send (B, C). The receiver reconstructs $K = B^a$ and removes the mask.

Under DDH, replacing K by a uniform $U \in \mathbb{G}$ produces a group one-time pad. We make this replacement precise in the security proof.

[Maj26, p. 5]; [HPS14, §2.4, pp. 70–73]; [BS23, Ex. 11.5, p. 459]; [ElG85].

Multiplicative ElGamal: the algorithms

Fix public parameters $(\mathbb{G}, q, g)$ and message space $\mathcal{M} = \mathbb{G}$.

Key generation

$a \xleftarrow{\$} \mathbb{Z}_q$, $A = g^a$. Return $\mathsf{pk} = A$ and $\mathsf{sk} = a$.

Encryption $\mathsf{Enc}_A(m)$

Sample fresh $b \xleftarrow{\$} \mathbb{Z}_q$.

$$\boxed{(B, C) = (g^b, mA^b).}$$

The ciphertext belongs to $\mathbb{G}^2$.

Decryption $\mathsf{Dec}_a(B, C)$

Require $B, C \in \mathbb{G}$, then return

$$\boxed{m = C\,(B^a)^{-1}.}$$

Division means group inversion.

Reject invalid encodings and elements outside $\mathbb{G}$.

[Maj26, p. 5]; [HPS14, Table 2.3, p. 72]; [BS23, Ex. 11.5, p. 459].

ElGamal correctness and the same worked example

Proof idea. The sender's mask A^b equals the receiver's mask B^a. Cancel it in the group.

$$\mathrm{Dec}_a(g^b, mA^b) = mA^b\big((g^b)^a\big)^{-1} = mg^{ab}(g^{ab})^{-1} = m.$$

Encrypt $m = 9$ modulo 23

Use $A = 8$ and fresh $b = 7$.

$$B = 13, \quad K = 12, \quad C = 9 \cdot 12 \equiv 16.$$

The ciphertext is $(13, 16)$.

Decrypt with $a = 3$

$$B^a = 13^3 \equiv 12, \qquad 12^{-1} \equiv 2.$$

Hence

$$\mathrm{Dec}_3(13, 16) = 16 \cdot 2 \equiv 9.$$

[HPS14, Table 2.3, p. 72]; [BS23, Ex. 11.5, p. 459]; arithmetic checked independently.

The privacy goal: a chosen-message challenge

Public-key IND-CPA guessing game

1. Challenger generates (A, a) and gives the public key A to attacker $\mathcal{A}$.
2. $\mathcal{A}$ chooses two messages $m_0, m_1 \in \mathbb{G}$ and retains its state.
3. Challenger samples $s \xleftarrow{\$} \{0,1\}$ and returns $c^* \xleftarrow{\$} \mathsf{Enc}_A(m_s)$.
4. $\mathcal{A}$ outputs $\widehat{s}$ and wins if $\widehat{s} = s$.

$$\boxed{\mathsf{Adv}^{\mathrm{guess}}_{\mathrm{CPA}}(\mathcal{A}) = \left|\Pr[\widehat{s} = s] - \tfrac{1}{2}\right|.}$$

The attacker can encrypt any messages using the public key. A separate encryption oracle gives no additional power. The game gives no decryption oracle.

[BS23, Attack Game 11.1, pp. 436–437; Def. 11.2, p. 437].

The reduction: turn an encryption attack into a DDH test

Proof idea. The challenge supplies a public key, ephemeral share, and mask.

On challenge (X, Y, T), distinguisher D runs:

1. Give $A = X$ to $\mathcal{A}$, obtaining $m_0, m_1 \in \mathbb{G}$ and its state.
2. Sample $s \xleftarrow{\$} \{0,1\}$ and give $(Y, m_s T)$ as the challenge ciphertext.
3. Obtain $\widehat{s}$ from $\mathcal{A}$. Output 1 iff $\widehat{s} = s$.

If $T = g^{ab}$

$X = g^a$, $Y = g^b$, so

$$(Y, m_s T) = \mathrm{Enc}_X(m_s).$$

The simulation is exact.

If $T = U \xleftarrow{\$} \mathbb{G}$

The ciphertext is $(g^b, m_s U)$. The mask is fresh and independent, even conditional on the public key and chosen messages.

Reduction proving [BS23, Ex. 11.5(a), p. 459].

The random-mask experiment hides the message perfectly

Condition on the public key, pre-challenge state, messages, and Y. For $c \in \mathbb{G}$ and $s \in \{0,1\}$,

$$\Pr[m_s U = c] = \Pr[U = m_s^{-1} c] = \frac{1}{q}.$$

A group one-time pad

Multiplication by m_s permutes $\mathbb{G}$. The ciphertext distribution is independent of s, so

$$\Pr[D(\mathcal{D}_1) = 1] = \Pr[\widehat{s} = s] = \tfrac{1}{2}.$$

The exact reduction with our conventions

$$\boxed{\mathrm{Adv}_{\mathrm{DDH}}(D) = \mathrm{Adv}^{\mathrm{guess}}_{\mathrm{CPA}}(\mathcal{A}).}$$

DDH makes this negligible. Difference-based CPA advantage is twice our guessing advantage.

Proof of [BS23, Ex. 11.5(a), p. 459]; CPA difference convention: [BS23, eq. (11.1), p. 437].

The message space and the fresh exponent matter

Messages must lie in the mask's group

If $\mathbb{G} \subsetneq \mathbb{F}_p^*$ but $m \in \mathbb{F}_p^*$ is unrestricted, then $C = mA^b \in m\mathbb{G}$: the ciphertext reveals the message's coset. For $\mathbb{G} = \mathrm{QR}_p$, the Legendre symbol of C equals that of m.

A repeated exponent repeats the mask

If two ciphertexts reuse b, then

$$C_1 = m_1 A^b, \quad C_2 = m_2 A^b \qquad \Longrightarrow \qquad C_1 C_2^{-1} = m_1 m_2^{-1}.$$

A known plaintext m_1 then reveals $m_2 = m_1 C_2 C_1^{-1}$.

For $t \in \mathbb{G}$, ElGamal is also malleable: (B, tC) decrypts to tm. IND-CPA privacy allows this behavior; it does not provide ciphertext integrity or CCA security.

[BS23, Ex. 11.5(c), Ex. 11.6, p. 459; §16.1.4, p. 688].

Candidate groups for the DDH assumption

A prime-order subgroup of $\mathbb{F}_p^*$

Choose primes $p = kq + 1$ and q, and use the subgroup of k-th powers. It has order q. For a safe prime $p = 2q + 1$, this is QR_p.

Curve groups

Suitable large prime-order subgroups of elliptic curves, and suitable subgroups of hyperelliptic Jacobians, are additional candidates. The group family and its parameters are part of the assumption.

DDH hardness is a conjecture for a specified family. Prime order by itself does not establish it: some curve groups admit efficient pairing-based DDH tests.

The safe-prime quotient $\mathbb{F}_p^*/\{\pm 1\} \cong \mathrm{QR}_p$ appears in the appendix.

[Maj26, p. 6]; [Bon98, §1.1, preprint p. 2]; [HPS14, Ex. 6.40, p. 369].

In the full multiplicative group, parity is visible

Now take an odd prime p and a generator g of the *full* group $\mathbb{F}_p^*$, whose order $p-1$ is even.

The Legendre symbol exposes exponent parity

For $X = g^x$, Euler's criterion gives

$$\chi(X) := X^{(p-1)/2} \bmod p = \begin{cases} +1 & X \text{ is a square,} \\ -1 & X \text{ is a nonsquare,} \end{cases} \qquad \boxed{\chi(g^x) = (-1)^x.}$$

Repeated squaring computes $\chi(X)$ efficiently.

Proof idea. A generator is a nonsquare, so $g^{(p-1)/2} = -1$. Raise this identity to the power x.

Exponent parity is well-defined modulo $p-1$. Recovering it does not require recovering the entire discrete logarithm.

[HPS14, Prop. 3.61, pp. 170–171; Example 3.69, eq. (3.35), p. 176]; [BS23, Lemma 16.1, p. 687].

The source's parity attack: compute the advantage

Use only the final element T. Output "real" if $\chi(T) = +1$ and "random" otherwise.

Real tuple: $T = g^{ab}$

ab is odd exactly when both a and b are odd. Hence

$$\Pr[\chi(T) = +1] = 1 - \tfrac{1}{2} \cdot \tfrac{1}{2} = \tfrac{3}{4}.$$

Random tuple: $T = g^{r}$

Exactly half the exponents in $\mathbb{Z}_{p-1}$ are even. Hence

$$\Pr[\chi(T) = +1] = \tfrac{1}{2}.$$

$$\boxed{\mathrm{Adv}_{\mathrm{DDH}} = \tfrac{1}{4}, \qquad \Pr[\widehat{u} = u] = \tfrac{5}{8}, \qquad \text{guessing advantage} = \tfrac{1}{8}.}$$

Thus DDH fails in the full group $\mathbb{F}_p^*$. In an odd prime-order quadratic-residue subgroup, every element has Legendre symbol $+1$, so this test gives no information.

Completes the calculation in [Maj26, p. 9]; [Bon98, §1.1, preprint p. 2].

DL, CDH, and DDH ask different questions

Problem	Input	Required output
DL	$X = g^a$	The exponent $a \in \mathbb{Z}_q$
CDH	$X = g^a$, $Y = g^b$	The group element g^{ab}
DDH	$X = g^a$, $Y = g^b$, T	Distinguish $T = g^{ab}$ from independent uniform T

An efficient algorithm for DL gives one for CDH

Recover $a = \log_g X$, then output $Y^a = g^{ab}$.

An efficient algorithm for CDH gives a DDH test

Compute $Z = \mathrm{CDH}(X, Y)$ and output "real" iff $T = Z$.

[Maj26, pp. 8, 10]; [HPS14, p. 69; Ex. 2.7, p. 108]; [BS23, §10.5, pp. 409–411].

The reduction's exact advantage

Normalize a CDH algorithm C to output an element of $\mathbb{G}$, replacing any invalid output by 1. Let ρ be the normalized algorithm's success probability on independent uniform $a, b \in \mathbb{Z}_q$.

Define $D(X, Y, T) = [C(X, Y) = T]$.

Real challenge

$$\Pr[D(\mathcal{D}_0) = 1] = \rho.$$

Random challenge

Given (X, Y) and C's coins, its output is fixed and T is independent uniform.

$$\Pr[D(\mathcal{D}_1) = 1] = \tfrac{1}{q}.$$

$$\boxed{\mathrm{Adv}_{\mathrm{DDH}}(D) = \left|\rho - \tfrac{1}{q}\right|.}$$

An exact CDH solver gives gap $1 - 1/q$ and guessing success $1 - 1/(2q)$. An exact DL solver gives the same test.

Quantifies the source's reductions [Maj26, pp. 8, 10]; [BS23, §10.5, p. 411].

What the implication arrows mean

Algorithms: stronger computational ability gives weaker tasks

$$\text{DL solver} \implies \text{CDH solver} \implies \text{DDH distinguisher.}$$

Hardness: take the contrapositives

$$\boxed{\text{DDH hard} \implies \text{CDH hard} \implies \text{DL hard.}}$$

For growing groups, $1/q$ must be negligible; this is implicit in cryptographic parameter choices.

The reverse hardness implications do not follow from these reductions. In particular, the Legendre-symbol attack makes DDH easy in $\mathbb{F}_p^*$ without giving a general DL or CDH algorithm.

Return to the motivating question. An inability to recover an entire secret does not by itself establish that the secret looks uniformly random.

[Maj26, pp. 8–10]; [BS23, §10.5, p. 411; §16.1.3, p. 687]; [Bon98, §1, preprint pp. 1–2].

The lecture in three equations

The shared value is pseudorandom under DDH

$$(g^a, g^b, g^{ab}) \approx_c (g^a, g^b, U), \qquad U \xleftarrow{\$} \mathbb{G}.$$

ElGamal masks a group message with that shared value

$$\mathrm{Enc}_A(m) = (g^b, mA^b), \qquad \mathrm{Dec}_a(B, C) = C(B^a)^{-1}.$$

Replacing A^b by U gives a uniform second component.

The group can leak information even when DL is difficult

$$\chi(g^x) = (-1)^x \qquad \text{in the full group } \mathbb{F}_p^*.$$

The DDH assumption concerns the chosen group family, not exponentiation in every group.

[BS23, Def. 10.8, p. 411; Ex. 11.5, p. 459; Lemma 16.1, p. 687].

Check your understanding

1. What is the difference between CDH hardness and DDH hardness?
2. Which exponent is fixed with an ElGamal key, and which is fresh for every encryption?
3. Where does the ElGamal privacy proof use $m \in \mathbb{G}$?
4. Does the Diffie–Hellman identity prevent a man-in-the-middle attack?
5. If a DDH test has distinguishing advantage $1/4$, what is its oriented guessing success?

Answers. (1) Computing the shared element versus distinguishing it from uniform. (2) a is fixed; b is fresh. (3) $U \mapsto mU$ must permute the mask's group. (4) No; authentication is additional. (5) $1/2 + (1/4)/2 = 5/8$.

[BS23, §10.5, pp. 409–411; Ex. 11.5–11.6, p. 459]; [HPS14, Example 3.13, pp. 126–127].

Reading guide

Hoffstein–Pipher–Silverman, second edition [HPS14]

- §2.2–2.4, pp. 64–73: DL, Diffie–Hellman, and ElGamal.
- Tables 2.2–2.3, pp. 68, 72: protocol and encryption algorithms.
- Example 3.69, pp. 176–177, especially eq. (3.35): exponent parity.

Boneh–Shoup, version 0.6 [BS23]

- §10.4–10.5, pp. 406–411: key exchange and hardness games.
- Attack Game 11.1, pp. 436–437: public-key privacy.
- Exercises 11.5–11.6, p. 459: *multiplicative* ElGamal and its message space.
- §16.1.3, Lemmas 16.1–16.2, p. 687: even-order DDH attacks.

All textbook locators use printed page numbers. The main construction in Boneh–Shoup §11.5 uses hashing and symmetric encryption; Exercise 11.5 matches this lecture's construction exactly.

References I

[Bon98] Dan Boneh.

The decision Diffie–Hellman problem.

In Joe P. Buhler, editor, *Algorithmic Number Theory, Third International Symposium, ANTS-III*, volume 1423 of *Lecture Notes in Computer Science*, pages 48–63. Springer, 1998.

Slide locators refer to the author-hosted preprint.

URL: https://crypto.stanford.edu/~dabo/papers/DDH.pdf, doi:10.1007/BFb0054851.

[BS23] Dan Boneh and Victor Shoup.

A graduate course in applied cryptography, 2023.

Version 0.6, January 2023. Slide locators use printed page numbers.

URL: https://toc.cryptobook.us/book.pdf.

[DH76] Whitfield Diffie and Martin E. Hellman.

New directions in cryptography.

IEEE Transactions on Information Theory, 22(6):644–654, 1976.

doi:10.1109/TIT.1976.1055638.

References II

[EIG85] Taher ElGamal.
A public key cryptosystem and a signature scheme based on discrete logarithms.
IEEE Transactions on Information Theory, 31(4):469–472, 1985.
doi:10.1109/TIT.1985.1057074.

[HPS14] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman.
***An Introduction to Mathematical Cryptography*.**
Undergraduate Texts in Mathematics. Springer, New York, second edition, 2014.
doi:10.1007/978-1-4939-1711-2.

[Maj26] Hemanta K. Maji.
CS355, Lecture 25: DDH Assumption, Key Agreement, and ElGamal Encryption.
Instructor's lecture slides, Purdue University, 2026.
Supplied source file 25(1).pdf, 10 slides.

Appendix: a stronger parity test uses the whole tuple

For a primitive root g in $\mathbb{F}_p^*$, write

$$\epsilon(X) = \begin{cases} 0 & \chi(X) = +1, \\ 1 & \chi(X) = -1. \end{cases}$$

Output "real" iff $\epsilon(T) = \epsilon(A)\epsilon(B)$ in $\mathbb{Z}_2$.

Real and random acceptance probabilities

If $(A, B, T) = (g^a, g^b, g^{ab})$, parity multiplication always agrees:

$$\Pr[D(\mathcal{D}_0) = 1] = 1.$$

If T is independent uniform, its parity is an independent fair bit:

$$\Pr[D(\mathcal{D}_1) = 1] = \tfrac{1}{2}.$$

Thus the distinguishing advantage is $1/2$, guessing success is $3/4$, and guessing advantage is $1/4$. Using A and B improves the source's final-component-only test.

[BS23, Lemmas 16.1–16.2, p. 687]; [HPS14, eq. (3.35), p. 176].

Appendix: the safe-prime quotient is the same group

Let $p = 2q + 1$ be prime, with q an odd prime. Consider

$$\phi : \mathbb{F}_p^* \longrightarrow \mathrm{QR}_p, \qquad \phi(x) = x^2.$$

Proof idea. Squaring is a surjective homomorphism whose kernel consists of 1 and -1. Passing to cosets removes exactly this ambiguity.

An explicit isomorphism

$$\boxed{\overline{\phi} : \mathbb{F}_p^*/\{\pm 1\} \longrightarrow \mathrm{QR}_p, \qquad [x] \longmapsto x^2.}$$

The map is well-defined because $(-x)^2 = x^2$, and it respects multiplication. If $x^2 = y^2$, then $(x - y)(x + y) = 0$ in $\mathbb{F}_p$, so $[x] = [y]$. Every square has a preimage. Thus the map is bijective.

This identifies the two safe-prime candidate groups in the source; both have order q.

Source alternatives: [Maj26, p. 6]; square classification: [HPS14, Prop. 3.61, pp. 170–171]; [Bon98, §1.1].

Cyclic-group arithmetic

If g has order q, then $x \mapsto g^x$ is a bijection $\mathbb{Z}_q \to \mathbb{G}$. Exponentiation satisfies $(g^a)^b = g^{ab}$, and multiplication by any fixed $m \in \mathbb{G}$ permutes $\mathbb{G}$.

Sources: HPS, §2.2, p. 67; Tables 2.2–2.3, pp. 68, 72.

Quadratic characters in the full group

For odd prime p, primitive root g, and $X \in \mathbb{F}_p^*$,

$$\chi(X) = X^{(p-1)/2} \in \{+1, -1\}, \qquad \chi(g^x) = (-1)^x.$$

The first value distinguishes squares from nonsquares. Exponentiation computes it in polynomial time in $\log p$.

Precise sources: HPS, Prop. 3.61, pp. 170–171; Example 3.69, eq. (3.35), p. 176. Boneh–Shoup, Lemma 16.1, p. 687; Appendix A.2.3, items 4–8, pp. 1097–1098.

[HPS14, BS23].

Technical input: the computational assumption

DDH assumption for the specified group family

Let $\mathrm{Grp}(1^\lambda)$ specify efficiently represented cyclic groups $(\mathbb{G}_\lambda, q_\lambda, g_\lambda)$, with efficient uniform exponent sampling and negligible $1/q_\lambda$. For every probabilistic polynomial-time D,

$$\left|\Pr[D(g^a, g^b, g^{ab}) = 1] - \Pr[D(g^a, g^b, g^r) = 1]\right| \le \mathrm{negl}(\lambda),$$

Here $a, b, r \xleftarrow{\$} \mathbb{Z}_{q_\lambda}$ are independent and D receives the group description. Probabilities include all coins, including group generation when randomized.

Exact source. Boneh–Shoup, Attack Game 10.6, p. 410, and Definition 10.8, p. 411. Our passive key theorem follows immediately; our ElGamal reduction proves the claim posed in Exercise 11.5(a), p. 459.

DDH is the only computational hardness assumption in our privacy proofs. The algebraic identities and attacks require no hardness assumption.

[BS23, Attack Game 10.6, pp. 410–411; Ex. 11.5, p. 459].