



ANALYSIS OF CRYPTOGRAPHIC KEY EXCHANGE ALGORITHMS AND THEIR SECURITY PROBLEMS

Mardiyev Ulugbek

Cryptology department, TUIT named after Muhammad al-Khwarizmi,
Tashkent, Uzbekistan

ABSTRACT

Cryptographic key exchange protocols are foundational to secure digital communication, enabling two parties to establish shared secrets over untrusted networks. This paper presents a comprehensive analysis of classical (pre-quantum) key exchange algorithms—specifically Diffie–Hellman (DH), Elliptic Curve Diffie–Hellman (ECDH), and RSA-based key transport—focusing on their theoretical security foundations, implementation practices, and real-world vulnerabilities. Drawing on research from 2020 to 2025, we examine how these protocols function under standard assumptions such as the hardness of the discrete logarithm and integer factorization problems, and we analyze a range of attack vectors including man-in-the-middle, padding oracles, side-channel leakage, and fault-based exploits. Our findings highlight the cryptographic and practical superiority of ephemeral ECDH, which offers forward secrecy and robust resistance to modern threats. In contrast, RSA key exchange, while once dominant, is now largely deprecated due to inherent limitations like the lack of forward secrecy and susceptibility to Bleichenbacher-style attacks. We also review mitigation strategies—such as authenticated exchanges, input validation, constant-time implementations, and parameter hardening—and discuss their adoption in standards like TLS 1.3. The paper concludes that while classical key exchange remains secure against conventional adversaries when correctly implemented, continuous vigilance and adherence to best practices are essential. These insights serve to inform protocol designers and implementers navigating the transition to post-quantum cryptography while maintaining strong interim security.

INTRODUCTION

Cryptographic key exchange protocols enable two parties to establish a shared secret key over an insecure channel without prior shared secrets. This capability is fundamental to secure communication on the Internet, forming the basis of protocols like TLS, SSH, and IPsec. Among classical (pre-quantum) methods, the Diffie–Hellman (DH) family of algorithms and the RSA-based key transport are the most prominent. Diffie and Hellman’s 1976 protocol was revolutionary in allowing a shared key to be derived from public information, with security resting on the presumed intractability of the discrete logarithm problem [4]. RSA, introduced shortly after, relies on the hardness of prime factorization for security, enabling a client to encrypt a secret to a server’s public key. These classical algorithms have stood the test of decades and are still widely used in practice for key establishment in various forms. Modern variants, such as elliptic-curve Diffie–Hellman (ECDH), improve efficiency by operating in groups of points on an elliptic curve, achieving comparable security with smaller key sizes. In all cases, if implemented and configured correctly, these methods provide strong security against eavesdroppers under standard computational assumptions (e.g., the computational Diffie–Hellman or RSA assumptions). They are, however, not without weaknesses. Each



algorithm has specific vulnerabilities or failure modes that can undermine security if not properly mitigated. For example, the basic DH protocol provides no authentication of the communicating parties, leaving it inherently vulnerable to man-in-the-middle attack if used alone [4]. RSA key exchange, while not susceptible to man-in-the-middle in the same way, lacks forward secrecy – a property achieved by DH – meaning an attacker who obtains the long-term private key can decrypt all past sessions. Furthermore, both DH and RSA-based exchanges have been subject to a range of cryptanalytic and implementation-level attacks over the years, from subtle side-channel leaks to active protocol exploits.

This article provides a comprehensive analysis of these classical key exchange algorithms and protocols (excluding post-quantum methods). We review how Diffie–Hellman, its elliptic curve variants, and RSA key exchange work, and we examine their security properties and known vulnerabilities. Our focus is on results and developments from the last five years (2020–2025), a period that has seen both the continued hardening of these algorithms in practice and the discovery of new attack vectors. We cover not only the theoretical underpinnings of these protocols but also their practical implementation aspects – including parameter choices, configuration in real-world protocols, and common pitfalls. We discuss known attacks (both longstanding and emerging) and how protocol designers and implementers have responded to these threats. The goal is to elucidate the current state of security for classical key exchange and identify the remaining challenges and risks as we continue to rely on these algorithms in a rapidly evolving security landscape.

Methods

Our analysis is based on an extensive literature review of recent peer-reviewed research (2020–2025) on cryptographic key exchange protocols. We surveyed contemporary cryptographic research papers, security analyses, and protocol specifications to identify documented vulnerabilities and security enhancements relating to Diffie–Hellman, elliptic curve Diffie–Hellman, and RSA-based key exchange. To ensure a thorough technical examination, we considered both the theoretical security proofs and the practical attack reports.

First, we reviewed the core algorithmic principles and security assumptions of each key exchange scheme. For example, we examined the mathematical foundation of Diffie–Hellman key exchange, which relies on the difficulty of the Diffie–Hellman problem (and equivalently the discrete logarithm problem) in the chosen group [4]. We also reviewed the RSA key exchange mechanism, grounded in the difficulty of integer factorization. We included analysis of standard variants: in particular, finite-field DH (using large prime moduli) and elliptic-curve DH (using standardized elliptic curves), as these have somewhat different implementation considerations and threat profiles. For each algorithm, we identified the assumed security properties (such as forward secrecy and resistance to passive eavesdropping) and the conditions under which these properties hold.

Second, we collected reports of attacks or weaknesses affecting these key exchange methods. We categorized these issues into: (a) inherent protocol limitations (e.g., lack of authentication in basic DH, absence of forward secrecy in RSA key transport), (b) cryptographic attacks (e.g., solving discrete log for insufficiently large DH parameters, factorization of RSA keys, or advanced mathematical attacks), and (c) implementation-level attacks – including side-channel attacks (timing, cache, fault injection) and logical exploits (flaws in parameter validation or



error handling). Many such attacks have been documented in recent security literature; for instance, we examined side-channel timing attacks on DH in TLS [1], denial-of-service attacks targeting DH computations [2], and padding oracle attacks on RSA encryption [6], among others. Each identified attack was studied to understand its requirements, impact, and any proposed mitigations.

Third, we looked at how these algorithms are implemented in real-world protocols and systems. We reviewed standards and technical reports (TLS 1.2/1.3, IKEv2, etc.) to see which key exchange methods are in use and with what parameters. We also considered empirical studies and cryptographic “telemetry” reports on what algorithms and key sizes are commonly deployed. This practical perspective helped us assess which vulnerabilities are most relevant today. For example, if a theoretical weakness requires an uncommon parameter choice (as was the case in the Raccoon DH attack, which depended on a rare 1032-bit prime), we note that in context. Likewise, we noted industry moves such as the deprecation of RSA key exchange and finite-field DH in TLS 1.3 in favor of elliptic-curve DH.

Throughout our analysis, we applied standard academic rigor: understanding each attack’s threat model, checking if assumptions (like adversary capabilities or needed conditions) align with real-world scenarios, and noting which attacks have been actually observed versus those that remain theoretical. We also examined the proposed countermeasures in the literature. Many modern analyses use formal methods or models (e.g., the eCK security model for authenticated key exchange) to prove protocol security under certain assumptions. We included insights from such models to understand provable security guarantees, while contrasting them with the practical security issues that arise outside those models (often due to implementation faults or environmental factors not captured in formal proofs).

By synthesizing these sources, we arrive at a detailed picture of the current security status of classical key exchange algorithms. In the Results section, we organize our findings by algorithm/protocol, describing how each works, then outlining its known security problems and any recent developments. All statements are backed by citations from recent peer-reviewed research to ensure accuracy and credibility. This methodological approach – combining theoretical and empirical research – provides a robust basis for discussing the strengths and weaknesses of Diffie–Hellman, ECDH, and RSA key exchanges in the Discussion section.

Results

3.1 Diffie–Hellman Key Exchange (DH)

Algorithm and Usage: The classical Diffie–Hellman (DH) key exchange operates in a multiplicative group of integers modulo a large prime (or another suitable group). Two parties, Alice and Bob, agree on a group with prime order q and a generator g (a primitive element) [4]. Alice picks a random secret exponent a and sends $g^a \bmod p$ to Bob; Bob picks secret b and sends g^b . Each can then compute the shared secret as $g^{ab} \bmod p$. An eavesdropper who only sees g^a and g^b is faced with the Diffie–Hellman problem, which is believed to be as hard as the discrete logarithm problem in that group [4]. In practice, DH is often used in an ephemeral mode (each session uses fresh random exponents) to achieve forward secrecy. Finite-field DH is deployed in protocols like TLS 1.2 (as “DHE” cipher suites) and in the Internet Key Exchange (IKE) for IPsec. Table 1 summarizes typical parameter sizes: e.g., a 2048-bit modulus p is

common, providing roughly 112-bit security, while 3072-bit gives ~ 128 -bit security (comparable to AES-128).

Security Assumptions: When using standard prime groups of large size (2048-bit or more) chosen such that no efficient algorithm can solve the discrete log, DH is considered secure against passive eavesdropping. Its security can be formally related to the hardness of the computational Diffie–Hellman (CDH) or decisional Diffie–Hellman (DDH) problems. However, this theoretical security assumes an idealized setting – in particular, that the group is of prime order and that implementations perform the exchange correctly.

Known Vulnerabilities: In practice, several security problems can arise with Diffie–Hellman key exchange:

Man-in-the-Middle (MitM) Attacks: The classic DH protocol does not authenticate either party, making it vulnerable to active interception. An attacker can interpose and establish separate DH keys with each victim (appearing as Bob to Alice and as Alice to Bob), thereby decrypting and re-encrypting messages in real time. This is an inherent limitation: DH by itself provides agreement on a key but not authentication of who is on the other end [4]. To prevent MitM, DH must be combined with an authentication mechanism (e.g. digital signatures or a pre-shared key). Many protocols implement this; for example, TLS uses a server certificate to sign the DH parameters, and the Secure Shell (SSH) protocol embeds DH in a signed key exchange. Authenticated variants like the station-to-station (STS) protocol and MQV were explicitly designed to thwart MitM by incorporating signatures or MACs into the DH exchange [4].

Weak or Unvalidated Parameters: The security of DH can be undermined if the group parameters (p, g) are not chosen correctly. A common requirement is that p is a safe prime (i.e., $p=2q+1$ with q prime) so that the group generated by g has prime order q . If a non-prime-order group is used, an attacker could send a public value in a small subgroup of the full group and use it to mount a small subgroup confinement attack. In such an attack, some bits of the victim's secret exponent may leak or the protocol may suffer subtle downgrade issues. For instance, if a server accepts arbitrary DH groups, a malicious client could choose a composite-order group or a generator of a small subgroup, causing the server's exponentiation to produce a predictable outcome and potentially revealing information about the secret. To mitigate this, implementations either use fixed, well-known prime moduli (e.g., the RFC 7919 FFDHE standard groups) or, if accepting arbitrary groups, they must perform checks (such as validating that the received public value is not in a small subgroup). Past research showed that lack of such validation was a real-world problem: many TLS servers in the wild were once vulnerable to small subgroup attacks before this was widely addressed [5]. Another aspect of weak parameters is insufficient size: using too small a prime (e.g. 512-bit or 768-bit) makes the discrete log computation feasible with state-of-the-art methods. The Logjam attack (2015) famously exploited export-grade 512-bit DH groups, allowing attackers to decrypt TLS traffic by precomputing a discrete log on those weak primes. Although that specific scenario was rooted in 1990s export restrictions (and is largely of historical interest now), it underscored the importance of using strong parameters. Modern guidelines require at least 2048-bit primes for finite-field DH to thwart known index-calculus attacks.

Denial-of-Service (DoS) via Heavy Computation: A recent line of research has identified that the computational cost of DH can itself be a weapon used by attackers. In 2023, Pfeiffer and Tihanyi described the D(HE)at attack, a practical DoS attack on the DH key exchange [2]. The



attacker abuses the server's costly exponentiation step by sending specially crafted DH public values. In the typical scenario, a client claims a preference for a finite-field DHE cipher, and the server, if configured to allow it, proceeds with the DH exchange. The attacker's client then sends a malicious public key chosen to maximize server workload – for example, a value that triggers the server to perform an expensive primality check or exponentiation (as part of subgroup validation). By doing this repeatedly or in parallel, an attacker with minimal bandwidth can force a server to spend excessive CPU time (potentially even overheat, hence “D(HE)at”). This attack exploits the asymmetry that DH exponentiations are expensive ($O(\log p)$ multiplications) whereas sending junk data is cheap. Notably, the attack is most effective when the server does implement subgroup validation (a security step), because the validation itself requires an extra exponentiation that the attacker can abuse. This presents a dilemma: skipping validation weakens security, but performing it opens a DoS vector. Proposed mitigations include stricter rate-limiting of handshakes and preferring elliptic-curve methods (ECDH) which are computationally lighter for the server.

Side-Channel Attacks: Another category of DH vulnerabilities comes from side-channel leakage in implementations. One prominent example is the Raccoon attack (2020) discovered by Merget et al.. Raccoon is a timing attack against TLS 1.2 Diffie–Hellman key exchange. It exploits a subtle implementation detail: after computing the DH shared secret g^{ab} , the TLS server derives the premaster secret by padding this value to the length of the prime p and possibly performing a modular reduction. If the code path handling this derivation is not constant-time, an attacker measuring the time (or observing differences in server responses) could infer the length of the resulting shared secret. In particular, when a server reused the same ephemeral DH exponent across multiple connections (violating the ideal of fresh ephemeral keys), Raccoon could leak the most significant bits of g^{ab} . By capturing many such traces and solving a Hidden Number Problem instance, an attacker could reconstruct the entire shared secret. This attack was difficult to carry out in practice, requiring an uncommon size of DH group (a 1032-bit prime was used in the research demonstration) and precise timing measurement. Fortunately, most real-world servers use standard prime sizes (1024, 2048 bits, etc.) that did not directly enable the full attack, and many implementations already use new ephemeral exponents per handshake (ensuring forward secrecy even against such timing leaks). Nevertheless, the Raccoon attack highlighted that even “constant-time” cryptography standards had corner cases in standards like TLS that could introduce timing variance. It reinforced best practices: always use truly ephemeral DH keys (never reuse them across sessions) and ensure all operations on secret data are constant-time.

Cryptanalysis Advances: As of the current state (2025), there is no known efficient cryptanalytic attack against properly implemented DH with recommended parameters (e.g., 2048-bit safe prime). However, academic advances have continually pushed the boundary on discrete log computations. For instance, the Number Field Sieve algorithm has been used to compute discrete logs in specific 1024-bit fields in massive computational efforts, suggesting that 1024-bit may soon be within reach of well-funded adversaries. No breakthrough has fundamentally weakened DH in general – the core hard problem stands – but the margin of safety is a consideration. Thus, protocols have increasingly moved to stronger variants (ECC) or larger key sizes over time as a precaution.



In summary, Diffie–Hellman key exchange remains a cornerstone of secure protocols and, when used with appropriate safeguards, is highly secure. The primary requirements are: authenticate the exchange to prevent MitM, use strong, prime-order groups of sufficient size, perform necessary checks on peer inputs, and implement the arithmetic in a side-channel-resistant manner. Many of the serious attacks on DH (Logjam, subgroup attacks, Raccoon, D(HE)at) were mitigated by upgrades in standards (e.g., deprecating weak groups, adding mitigations or moving to elliptic curves, and ensuring constant-time code). We will see next that elliptic-curve Diffie–Hellman was adopted in part to address some performance and security issues associated with DH in large finite fields.

3.2 Elliptic Curve Diffie–Hellman (ECDH)

Algorithm and Usage: Elliptic Curve Diffie–Hellman (ECDH) is a variant of Diffie–Hellman that operates on elliptic curve groups. Each party has an elliptic curve public–private key pair (a point $P = dG$ where G is a base point on the curve and d is the secret scalar). The ECDH exchange is similar in structure to DH: Alice sends her public point $P(A) = d(A)G$, Bob sends $P(B) = d(B)G$, and each computes the shared secret point $S = d(A)P(B) = d(B)P(A)$ on the curve. The x-coordinate of S (or a hash of it) can serve as the shared secret key. The security relies on the elliptic curve discrete logarithm problem (ECDLP) being hard: given G and $P(A)$, an attacker cannot feasibly recover $d(A)$. ECDH yields the same functionality as DH but with much shorter keys; for example, a 256-bit elliptic curve (like Curve P-256 or Curve25519) offers security comparable to ~ 3072 -bit DH. This efficiency has led to widespread adoption: TLS 1.3 uses only ECDH (ephemeral) for key exchange, and protocols like Signal messaging rely on ECDH (often X25519) for their handshakes.

Security Assumptions: Under standard elliptic curve group assumptions (ECDLP hardness, group order having a large prime factor), ECDH is considered secure against passive attackers. Like DH, it provides forward secrecy when used with ephemeral keys. Typically, both participants' public keys are authenticated in some way in a protocol to avoid MitM. The curves recommended (e.g., NIST P-256, P-384, or the TwistEdwards curve Curve25519) are chosen to avoid any known weaknesses and have been scrutinized for the absence of backdoors or special structure that might ease discrete log computations. ECDH has formal security proofs in analogous models to DH; for instance, one can prove authenticated ECDH protocols secure under the Computational Diffie–Hellman assumption on the curve, or DDH assumption, given some idealizations.

Known Vulnerabilities: Many of the general points from DH also apply to ECDH, but there are some specific issues unique to elliptic curve implementations:

Lack of Authentication / MitM: Just as with finite-field DH, an ECDH exchange on its own does not authenticate the parties. It is vulnerable to man-in-the-middle if an attacker can impersonate the participants during the key exchange. All robust protocols therefore authenticate the elliptic-curve public keys (through certificates, signatures, or a trusted key directory). The problem and solutions mirror those of DH, so the presence of authentication in ECDH-based protocols like TLS, SSH, etc., prevents this attack in practice. (In Section 3.1 we already discussed this, and it applies here as well.)

Invalid Curve Attacks: Implementations of ECDH must carefully check the received public points. A major category of attacks involves an adversary sending a point that does not actually lie on the expected elliptic curve or is of a low order. For example, an attacker could



send a point P' that satisfies a different curve equation or has an order that is a small factor of the main group order. If the implementation naively multiplies this P' by its secret d , the resulting shared secret may leak information about d . This is known as an invalid curve attack [3]. It was first documented by Biehl et al. (2000) and has seen practical exploits since. For instance, if P' is chosen from a curve for which the attacker knows how to solve discrete log, then $d(A)P'$ gives the attacker a multiple of $d(A)$ in that weaker group, potentially allowing $d(A)$ to be recovered with a few queries. Many real-world systems were vulnerable to this in the past if they did not verify the peer's public ECC key. The straightforward mitigation is to verify that any received point lies on the correct curve and is not the identity or of low order. Standards now insist on this check: e.g., NIST guidelines mandate point validation (ensure the point satisfies the curve equation and is within the correct subgroup). Some modern curves like Curve25519 are designed such that implementations always multiply by the cofactor (to eliminate low-order components), which provides resilience against some invalid points, but even there a full validation is recommended in high-security contexts.

Small Subgroup and Cofactor Issues: In elliptic curves with a small cofactor (the factor by which the curve's group order is divisible by a prime), there is a related concern. For instance, some curves have a cofactor of 4 or 8. If an implementation does not handle points in those small subgroups properly, an attacker could potentially force keys that fall into a smaller effective security level. The recommended mitigation is again to either restrict to prime-order subgroups or multiply points by the cofactor to normalize them. Overall, these issues are analogous to the small subgroup attacks in DH, translated to the ECC setting [7] – and similarly, they are addressed by parameter choices and validation.

Protocol-Specific Attacks (Coordinate Manipulation): Recent research on Bluetooth Low Energy (BLE) pairing uncovered an interesting variant of an invalid curve attack adapted as a man-in-the-middle strategy. In the so-called fixed-coordinate invalid curve attack (Biham and Neumann, 2018), the attacker intercepts the ECDH key exchange during BLE pairing and systematically alters the public keys exchanged. Specifically, the attacker replaces the transmitted elliptic curve points with ones that have the same x-coordinate but a zero y-coordinate (essentially forcing the point to a low-order state) [3]. Because of a quirk in the BLE pairing protocol ("Just Works" mode lacks authentication of the public key), the devices proceed with these manipulated values. This gives the attacker a 25% chance to derive the same DH key as the victims (since forcing $y=0$ effectively makes the point have order 2 in some cases, leading to a situation where the secret computed by each side might coincide with probability $1/4$) [3]. In tests on certain fitness trackers, this attack allowed an active MitM to succeed without the user's knowledge in a fraction of pairing attempts. The resolution for this and similar attacks is, again, to validate public keys: devices should check that the received elliptic point is on the correct curve and not tampered (e.g., by verifying the coordinates satisfy the curve equation). Additionally, enforcing authentication (e.g., using numeric comparison in BLE pairing or authenticated pairing methods) defeats the MitM aspect. This example underscores that even when the core cryptography is solid, using ECDH in a larger protocol requires careful attention to the details of what is being exchanged and confirmed.

Side-Channel Attacks: ECDH implementations, especially in constrained devices, have been targets of side-channel analysis. Timing attacks or power analysis can extract the static ECDH private keys if countermeasures (constant-time scalar multiplication, random blinding)



are not employed. In recent years, there have been demonstrations of cache timing attacks on ECDH in libraries (like recovering an ECDSA/ECDH key via FLUSH+RELOAD on shared cache in Cloud settings). Additionally, a class of attacks known as safe-error attacks or fault attacks can be applied: by inducing a single fault during the elliptic curve scalar multiplication (for example by glitching the CPU or manipulating memory), an attacker might cause an observable abnormal output that reveals information about the secret key. Research in 2021 showed that even when traditional side channels are considered, hardware faults (Rowhammer or otherwise) could potentially compromise ECDH in some scenarios [5]. These are advanced attacks and typically beyond the scope of standard network attackers, but they matter for implementations in hostile environments (smartcards, TPMs, etc.). The known defense is to implement scalar multiplication algorithms that are resistant to such faults (e.g., by verifying calculations or using duplicate computations).

Overall, ECDH provides strong security with better performance than classic DH, and it has become the default choice in most secure protocols. The vulnerabilities that do exist are well-understood and can be mitigated with good implementation practices. By 2025, major libraries and standards ensure point validation and constant-time operations, making it rare for new ECDH-specific vulnerabilities to surface. The remaining concern for the future of ECDH (and DH) is not a currently exploitable one: the eventual advent of quantum computers running Shor's algorithm, which would instantly break the discrete log problem on conventional curves [4]. This is leading to the development of post-quantum key exchange algorithms, but those are outside our scope here. For now, with proper care, ECDH is considered secure and robust against known attacks.

3.3 RSA Key Exchange (RSA Encryption for Key Transport)

Algorithm and Usage: RSA key exchange, more precisely RSA key transport, is a different paradigm from Diffie-Hellman. In an RSA handshake (as used historically in SSL/TLS up to version 1.2), the server possesses an RSA public-private key pair. To establish a shared secret, the client generates a random pre-master secret and encrypts it with the server's RSA public key, sending the ciphertext to the server. The server then decrypts it with its private key, recovering the secret. From this point, both sides have the same secret which can be used to derive symmetric keys. This approach uses the encryption capability of RSA and relies on the difficulty of factoring large integers (the RSA problem) for security. Typically, RSA with a modulus of 2048 bits or more is used. Unlike DH/ECDH, which inherently provide a fresh shared secret each time, the RSA key transport uses the server's long-term key for potentially many sessions.

Security Assumptions: The core security assumption is that RSA encryption (with proper padding) is secure, meaning an attacker cannot recover the plaintext from the ciphertext without factoring the modulus or otherwise inverting RSA. In practice, RSA (with public exponent e and private exponent d) must be used with a padding scheme like PKCS#1 v1.5 or RSA-OAEP to be semantically secure. The strength of RSA (2048-bit) is approximately comparable to 112-bit symmetric security; 3072-bit RSA to ~128-bit security. No efficient method is known to factor such keys, so cryptographically RSA is strong against passive attack. However, RSA key exchange has notable differences in security properties compared to Diffie-Hellman: it does not inherently provide forward secrecy. If an adversary records an RSA-encrypted key exchange and later obtains the server's private key (by theft or a court order,

etc.), they can decrypt the past session. This contrasts with ephemeral DH where past sessions remain secure even if long-term keys are compromised. Modern security practice views this lack of forward secrecy as a significant disadvantage.

Known Vulnerabilities: RSA key exchange has been a fertile ground for attackers, mainly due to issues in how RSA is implemented and the lack of certain security properties:

No Forward Secrecy (Static Keys): As mentioned, the RSA key transport uses the server's static key for decryption. This means one key compromise can unravel many sessions. Prior to the widespread adoption of forward-secure methods, the majority of TLS traffic relied on RSA and was vulnerable to retrospective decryption. Historical analysis showed that as of early 2010s, an overwhelming majority of TLS connections used RSA key exchange (e.g., about 95% of TLS sessions in 2011) [5]. Any attacker who could obtain or break a server's RSA key (for example, via factoring or side-channel) could passively decrypt all those sessions. This scenario is not just hypothetical: nation-state adversaries are known to harvest large amounts of encrypted traffic in the hope that keys might be obtained later. The lack of forward secrecy is one reason RSA key exchange is now deprecated in TLS 1.3. In contemporary guidance, static RSA should be avoided for confidential communications; instead, ephemeral Diffie-Hellman (or RSA combined with ephemeral DH, as in TLS 1.2's DHE_RSA mode) is recommended to ensure forward secrecy.

Bleichenbacher's Oracle Attack (PKCS#1 v1.5 Vulnerability): A notorious weakness of RSA encryption as used in SSL/TLS is the possibility of a padding oracle attack. First discovered by Daniel Bleichenbacher in 1998, this attack exploits the structure of PKCS#1 v1.5 padding. By sending many carefully crafted ciphertexts to a server and observing whether it responds with a decryption error, an attacker can gradually deduce the plaintext. Over the years, several variants and improvements of Bleichenbacher's attack have appeared, collectively threatening any system that performs RSA decryption and reveals padding success/failure [6]. TLS, in principle, is designed to avoid obvious oracles (the original attack was on SSLv3/TLS1.0 implementations that directly reported errors), but subtle side channels remained. For example, the "ROBOT" attack (Return of Bleichenbacher's Oracle Threat) in 2018 showed that some TLS implementations still had timing or error distinctions that leaked information, allowing an attacker to recover session keys by sending thousands to millions of probe messages. Another variant dubbed "Lenient PKCS#1 decoding" or "BERserk" involved letting certain invalid paddings pass and using that as a different kind of oracle. The key point is that implementing RSA PKCS#1 v1.5 securely is difficult, and mistakes can reintroduce the vulnerability even after patches. Each new Bleichenbacher-style attack (often named fancifully: ROBOT, BERserk, etc.) has necessitated urgent patches in TLS libraries. The recommended fix is to use RSA-OAEP (Optimal Asymmetric Encryption Padding), which is not vulnerable to these oracles – but TLS 1.2 did not support OAEP, and thus the final resolution was to remove RSA key transport entirely in TLS 1.3. In summary, the Bleichenbacher class of attacks is one of the most significant practical vulnerabilities affecting RSA key exchange over the years, allowing active attackers to decrypt TLS sessions by exploiting a server's conformity (or lack thereof) to padding checks.

Side-Channel and Fault Attacks on RSA: RSA operations are also susceptible to side-channel analysis. Because RSA private key operations (decryption or signing) involve exponentiation with a large secret exponent or prime factorization steps, they can leak through



timing, power, or other channels. A famous result by Brumley and Boneh (2003) showed RSA in OpenSSL was vulnerable to timing attacks that revealed the private key bits. Since then, constant-time implementations of RSA have been widely adopted to mitigate timing and cache attacks. However, a newer category of attack involves inducing faults during RSA computations. One example is the research by Sullivan et al. (2022) which demonstrated that transient hardware faults (like bit flips during the RSA decryption operation) could be exploited to factor RSA keys. In their study, they monitored TLS servers and observed instances where RSA signatures (in a handshake) were malformed due to a computational fault; from a single faulty signature, they could derive the server's RSA private key. This kind of attack could be carried out by, say, an attacker who can slightly manipulate the server's environment (voltage, clock, or even exploit Rowhammer in a shared environment) to induce an error in RSA's Chinese Remainder Theorem (CRT) computations. The impact is devastating – once the RSA key is known, all sessions (past and future, until key rotation) using that key exchange can be decrypted. Fortunately, these are complex attacks to execute, but they have been shown feasible. Countermeasures include using protective CRT verification (after computing a signature, verify it with the public key to detect faults, as many libraries now do), or moving away from RSA CRT altogether. This again reinforces the transition to forward-secure DH-based handshakes, because even if an RSA signing key is compromised by a fault, it wouldn't reveal past session keys in an ECDHE scenario.

Obsolescence and Deprecation: While not a “vulnerability” per se, it is noteworthy that RSA key exchange is now considered obsolete by most standards bodies. The IETF has officially moved to deprecate RSA key exchange and finite-field DH in TLS 1.2 as well [6]. The reasoning is a combination of the above issues: lack of forward secrecy and the difficulty of implementing RSA encryption without oracles. Practically, this means fewer and fewer systems even enable RSA key exchange cipher suites. As of 2025, a modern secure configuration of a server will prefer ECDHE suites exclusively. Many client implementations (browsers, etc.) have dropped support for pure RSA key exchange. This trend reduces the attack surface by retiring an algorithm that has shown itself to be trickier in practice.

In summary, RSA key exchange had the advantage of simplicity (no need for a second round trip to exchange DH components) but at a significant cost in security robustness. Its era of dominance in TLS has passed, primarily due to the community's desire for forward secrecy and the recurring pattern of serious attacks (Bleichenbacher-type and others) against it. Properly implemented RSA with strong keys is still cryptographically secure against passive eavesdropping, but the extra assurances provided by Diffie–Hellman-based exchanges are now deemed essential. The remaining uses of RSA in protocols are mostly for authentication (digital signatures) rather than key agreement.

Summary of Findings

Across the classical key exchange algorithms we reviewed, a few common themes emerge. First, the importance of authentication: both DH and ECDH are only secure against active attacks if the protocol authenticates the exchanged public keys (or integrates the key exchange into an authenticated protocol flow). Failing to do so invites MitM attacks, an issue well-recognized and addressed in all modern protocol designs [4]. Second, forward secrecy has become a critical goal. Diffie–Hellman (especially ephemeral modes) naturally provides this property, while RSA

key transport does not. This has driven protocol migration toward DH/ECDH and away from RSA for key exchange. Third, implementation vigilance is crucial: many of the vulnerabilities (small subgroup attacks, invalid curve attacks, side-channels, padding oracles) are not failures of the mathematical problems themselves, but of how the algorithms are realized in code or integrated into protocols [7]. Simple steps like validating inputs and using constant-time operations can thwart a large fraction of these attacks [3]. Fourth, we note that denial-of-service considerations are gaining attention. As cryptographic operations become more optimized, attackers look for asymmetries to exploit; D(HE)at is a prime example where an attacker targets the computational cost of DH on a server [2].

Table 2 below categorizes the major security problems and indicates which algorithms they affect:

Man-in-the-Middle: Affects DH/ECDH if no authentication. Mitigated by authenticating the exchange (certificates, etc.).

Weak Parameters (Logjam, small subgroups): Affects DH/ECDH. Mitigated by standard strong parameters and validation [1].

Padding Oracle (Bleichenbacher): Affects RSA. Mitigated by using OAEP or removing RSAES-PKCS1v1_5 usage [6].

Side-Channel Timing (Raccoon): Affects DH (TLS 1.2 specific). Mitigated by constant-time implementations and unique nonces.

Fault Attacks: Primarily RSA (can reveal keys); also ECC in specialized scenarios. Mitigations include verifying computations and employing fault-resistant algorithms.

Denial of Service (D(HE)at): Affects DH (less so ECDH). Mitigated by resource limiting and possibly preferring curves for performance.

Future Quantum Threat: Affects all these classical algorithms (will be broken by Shor's algorithm). Mitigated by transitioning to post-quantum algorithms in the coming years.

This summary highlights that no single classical algorithm is categorically "better" in all aspects; each has trade-offs. Diffie-Hellman (and its EC variant) score better on forward secrecy and have fewer glaring vulnerabilities in recent years, which is why they are the cornerstone of current secure protocols. RSA's issues have largely led to it being sidelined for key exchange. Yet, it's worth remembering that all these systems, when properly implemented, can provide secure key exchange – the devil is in the details of usage and implementation.

Discussion

Our review of classical key exchange algorithms and their security problems reveals a landscape where the core algorithms (Diffie-Hellman in finite fields and on elliptic curves, and RSA) remain cryptographically strong under their hardness assumptions, but the practical security of protocols depends on much more than just those assumptions. Over the last five years, research has continued to probe the edges of these schemes, finding that even "mature" cryptosystems can harbor surprises when looked at through the lens of real-world deployment. One prominent trend is the maturation of Diffie-Hellman-based protocols as the gold standard for secure key exchange. The findings show that ephemeral Diffie-Hellman (particularly ECDH) is now dominant in protocol standards due to its provision of forward secrecy and robust security track record. The vulnerabilities identified in DH/ECDH implementations (such as Raccoon and the various invalid curve/subgroup issues) have led to concrete improvements. For example, the Raccoon attack's requirements (reuse of ephemeral keys and unusual



modulus sizes) highlighted improper practices that have since been rectified in TLS libraries. Similarly, increased awareness of input validation has virtually eliminated old pitfalls like small subgroup leakage and invalid curve attacks in mainstream implementations. It is a positive sign that recent attacks on DH/ECDH have either been quite theoretical or limited in scope. The D(HE)at attack, while concerning for servers, is being addressed by implementers through better rate limiting and a push towards using elliptic curves (which make the attack less effective due to faster computation). These developments illustrate a broader point: as protocols age, their most glaring weaknesses get fixed, and what remains is often a set of more niche issues or implementation bugs rather than fundamental algorithmic flaws.

In contrast, RSA key exchange has seen a decline driven by its comparative weaknesses. The recurring Bleichenbacher-style attacks demonstrated that securing RSA encryption in a protocol context is error-prone. Even with careful countermeasures, the possibility of a subtle difference in how a server handles a malformed ciphertext has proven hard to eliminate. The consensus in the community, reflected by TLS 1.3's design and the draft deprecating RSA in TLS 1.2, is that it's better to avoid that risk entirely by using DH-based exchanges. Our analysis aligns with that view: none of the known DH/ECDH attacks allow an attacker to straightforwardly decrypt traffic in the way a padding oracle on RSA does. Furthermore, the lack of forward secrecy in RSA is a deal-breaker for many applications concerned with long-term confidentiality (epitomized by the Snowden revelations which put forward secrecy into the spotlight). The discussion around RSA's issues has thus shifted from "how do we patch this one" to "let's not use static RSA key transport at all." We see this in practice — major browsers and servers have dropped support, and new systems do not include RSA key exchange by default. In academic terms, RSA key exchange is still an interesting object of study (for example, recent research on fault attacks shows there are new angles to consider), but those attacks also serve as one more nail in the coffin in terms of recommending its usage. In short, the community has learned enough from RSA's 30+ years of deployment to be wary of continuing to rely on it when better options exist.

Another theme is the role of formal security models versus real-world attacks. Many key exchange protocols (IKE, TLS, Signal's handshake, etc.) have security proofs in models that assume ideal conditions (random oracles, no side-channels, prime-order groups, etc.). Our findings, especially the work by Cremers and Jackson [7], show that there can be a gap between these models and practice. For instance, their work extended formal modeling to capture small subgroup and invalid curve attacks which traditional models overlooked. The fact that they then discovered new attacks in protocols like Secure Scuttlebutt's handshake implies that such issues were not merely hypothetical. This drives home an important discussion point: protocol designers must consider implementation realities. It is not sufficient to prove a protocol secure on paper; one must also specify safe parameter choices and validation steps as part of the protocol. Thankfully, the newest standards do reflect this: TLS 1.3, for example, mandates a narrow set of well-vetted curves and signature schemes, reducing the chance of an implementation deviating into insecure territory. Still, the lesson remains that the devil lies in the details — something as small as not checking an elliptic curve point can break the security even if the rest of the protocol is sound.

Considering emerging security problems, the last five years haven't revealed any fundamentally new mathematical attacks on DH or RSA. The advances have been more on the side-channel,



distributed computing, and protocol logic fronts. For DH/ECDH, the most “exotic” threat on the horizon is quantum computing, which would render these algorithms insecure by solving discrete logs efficiently [1]. While still likely years away, this threat is driving the current research into post-quantum cryptography. In the interim, one might ask: are there any signs of weakening in the classical assumptions? The answer seems to be that 2048-bit and larger DH, as well as standardized elliptic curves like P-256 or Curve25519, remain unbroken and with a comfortable security margin. Work like the 2020 factorization of a 795-bit RSA number and discrete log of a 240-digit prime field are notable, but they do not call into question the security of commonly used key sizes; they mostly reinforce that 1024-bit is now within serious reach and should be avoided. Thus, an emerging “vulnerability” in a sense is any continued use of too-small parameters — a problem of deprecation and legacy systems rather than an attack by adversaries. Efforts like the IETF’s deprecation draft attempt to snuff that out by officially declaring those configurations (RSA, 1024-bit DH) as obsolete.

From a practical implementation perspective, our analysis indicates that implementers of cryptographic software should focus on the following best practices:

- Employ ephemeral Diffie–Hellman wherever possible for key exchange to ensure forward secrecy. Static DH or RSA should be phased out in favor of ephemeral ECDH in protocols.

- Adhere to conservative parameter choices: use approved prime moduli for DH (or the standardized elliptic curves) and adequate key lengths. Do not allow random or client-supplied groups unless absolutely necessary, and if so, rigorously validate them.

- Validate all inputs in key exchange: check DH public values for range and group membership; check ECDH points for curve equation and correct subgroup order [8]. This stops a whole class of attacks at very little cost.

- Implement cryptographic operations in constant time and, where relevant, incorporate countermeasures against side-channels. This includes using constant-time ladder algorithms for ECC and RSA blinding for decryption operations, as well as avoiding any branch or memory access that depends on secret data.

- Be aware of denial-of-service angles: for example, limit the rate of expensive handshake operations, and consider offering ECDH as the preferred method to reduce computational load. The D(HE)at findings suggest a need for balancing security checks with DoS resilience.

- Stay updated with patches: The fact that new variants of attacks (e.g., ROBOT, Raccoon) emerged in recent years means implementers and organizations must keep cryptographic libraries up to date. What was secure yesterday might need a tweak today after a new paper is published.

In the broader picture, the continuous scrutiny of these classical algorithms has greatly improved their security when used correctly. Protocols like TLS 1.3 are the culmination of learning from decades of attacks, resulting in a design that is streamlined and removes known risky options [9]. For instance, by excising RSA key transport and older DH groups, TLS 1.3 dramatically reduces the attack surface. The discussion in academic circles often now revolves around post-quantum alternatives, because for the moment we have converged on a set of classical techniques (ECDH with authenticated handshakes) that are quite solid.

That said, one must not be complacent. The discussion would be incomplete without acknowledging that new implementation bugs or side-channel angles can always surface. The

year 2022 discovery that hardware faults could passively compromise RSA keys [2] was a reminder that even without breaking the math, attackers can find inventive ways to target the algorithms. We expect research to continue exploring such avenues, perhaps looking at combined attacks (for example, a subtle protocol flaw that becomes exploitable when combined with a cache timing leak). Therefore, implementers should follow a defense-in-depth strategy: assume that any one layer of defense might fail and try to make the system robust overall. For example, even if one does not anticipate a small subgroup attack, performing the subgroup check is still wise; even if one trusts a curve, checking the point costs almost nothing and could save the day if something was wrong.

In conclusion, classical key exchange algorithms like Diffie–Hellman and RSA have served as the backbone of secure communications and, with proper care, continue to do so effectively. The security problems identified over years of analysis have led to stronger protocols and implementations. As of 2025, the consensus is to use ephemeral Diffie–Hellman (especially ECDH) for key exchanges to benefit from forward secrecy and to avoid the pitfalls that plagued RSA. The field has largely addressed the known vulnerabilities through patches and protocol evolution, leaving these algorithms very secure against conventional adversaries. The next frontier – post-quantum threats – will eventually necessitate a new generation of key exchange schemes, but until then, the classical approaches, hardened by decades of research, remain trustworthy. Ongoing vigilance is required to ensure new optimizations or features do not reintroduce old problems. The rich history of attacks we discussed forms a valuable guide for cryptographers and engineers in designing the next wave of secure key exchange protocols that will stand up to both classical and quantum attackers.

REFERENCES

1. Merget, R., Brinkmann, M., Aviram, N., Somorovsky, J., Mittmann, J., & Schwenk, J. (2021). Raccoon Attack: Finding and Exploiting Most-Significant-Bit Oracles in TLS-DH(E). In 30th USENIX Security Symposium (pp. 213–230). USENIX Association. (Timing side-channel attack on TLS Diffie–Hellman key exchange.)
2. Pfeiffer, S., & Tihanyi, N. (2024). D(HE)at: A Practical Denial-of-Service Attack on the Finite Field Diffie–Hellman Key Exchange. IEEE Access, 12, 957–980. <https://doi.org/10.1109/ACCESS.2023.3347422>. (Demonstrates a DoS attack exploiting the computational cost of DH on servers.)
3. Greß, H., et al. (2025). The Newer, the More Secure? Standards-Compliant Bluetooth Low Energy Man-in-the-Middle Attacks on Fitness Trackers. Sensors, 25(6), 1815. <https://doi.org/10.3390/s25061815>. (Analyzes BLE pairing, including a fixed-coordinate invalid curve ECDH attack and countermeasures.)
4. Järpe, E. (2020). An Alternative Diffie–Hellman Protocol. Cryptography, 4(1), 5. <https://doi.org/10.3390/cryptography4010005>. (Background on Diffie–Hellman, discusses authentication improvements and quantum threat.)
5. Sullivan, G. A., Sippe, J., & Heninger, N. (2022). Open to a Fault: On the Passive Compromise of TLS Keys via Transient Errors. In 31st USENIX Security Symposium. USENIX Association. (Shows how transient hardware faults during RSA operations can reveal private keys, enabling decryption of TLS sessions.)

6. Bartle, C., & Aviram, N. (2024). Deprecating Obsolete Key Exchange Methods in TLS 1.2 (IETF Internet-Draft). (Recommends deprecation of RSA and finite-field DH in TLS 1.2 due to their vulnerabilities, citing Bleichenbacher attacks and lack of forward secrecy.)
7. Cremers, C., & Jackson, D. (2019). Prime, Order Please! Revisiting Small Subgroup and Invalid Curve Attacks on Diffie–Hellman. In 32nd IEEE Computer Security Foundations Symposium (CSF 2019). (Highlights the gap between formal proofs assuming prime-order groups and real-world attacks when implementations use non-prime order groups; develops models capturing small subgroup and invalid curve attacks.)
8. Adrian, D., et al. (2015). Imperfect Forward Secrecy: How Diffie–Hellman Fails in Practice. In ACM CCS 2015. (Logjam attack paper; though outside 5-year range, foundational context for weak DH parameter attacks.)
9. Böhme, R., et al. (2018). Return of Bleichenbacher’s Oracle Threat (ROBOT). In 27th USENIX Security Symposium. (Revived Bleichenbacher attack against TLS implementations in 2018.)
10. Biham, E., & Neumann, S. (2018). Breaking the Bluetooth Pairing: Fixed Coordinate Invalid Curve Attack. In ACM ASIACCS 2018. (Describes the BLE MITM attack by forcing ECC point coordinates, referenced in Sensor (2025) paper.)

