Revisit of KD04

Xianhui Lu[1], Xuejia Lai[2], Dake He[1], Guomin Li [1]
Email:lu_xianhui@sohu.com
1:School of Information Science & Technology, SWJTU, Chengdu, China
2:Dept. of Computer Science and Engineering, SJTU, Shanghai, China

May 21, 2007

Abstract

KD04 proposed by K. Kurosawa and Y. Desmedt is the most efficient public key encryption scheme provably secure against adaptive chosen ciphertext attack in standard model based on decision diffie-hellman problem. We proposed a simplify version of KD04 which is more efficient than KD04 while still can be proved to be secure against adaptive chosen ciphertext attack in standard model based on decision diffie-hellman problem.

Keywords: hybrid encryption, public-key encryption, standard model, chosen cipher-text attack

1 Introduction

Security against adaptive chosen cipher-text attacks (CCA secure) [1, 2, 3] is a strong and very useful notion of security for public-key encryption schemes. This notion is known to suffice for many applications of encryption in the presence of active attackers, including secure communication, auctions, voting schemes, and many others. CCA secure is commonly accepted as the security notion of choice for encryption schemes that are to be plugged in to a protocol running in an arbitrary setting [4, 5]. The random oracle model is a useful tool in constructing CCA secure public-key encryption schemes, but it does not rule out all possible attacks [6]. Schemes that can be proven to be CCA-secure in the standard model (without the use of heuristics such as random oracles) is more practical. Cramer and Shoup showed the first provably secure practical public-key encryption scheme in the standard model [7]. It is CCA secure under the Decisional Diffie-Hellman(DDH) assumption. They further generalized their scheme to projective hash families [8].

Public key encryption schemes often limit the message space to a particular group, which can be restrictive when one wants to encrypt arbitrary messages. For this purpose hybrid schemes are devised. In these cryptosystems a symmetric encryption scheme is used to overcome the problems typically associated with encrypting long messages using "pure" asymmetric techniques. This is typically achieved by encrypting the message with a symmetric encryption scheme and a randomly generated symmetric key. This random symmetric key is then somehow encrypted using an asymmetric encryption scheme. This approach has been successfully used for many years.

One important advance in hybrid cryptography is the development of the KEM/DEM model for hybrid encryption algorithms [12]. This model splits a hybrid encryption scheme into two distinct

components: an asymmetric key encapsulation mechanism (KEM) and a symmetric data encapsulation mechanisms (DEM). In order to obtain a CCA-secure hybrid encryption, it is sufficient that both KEM and DEM are CCA-secure. (Accordingly, we refer the framework of [12, 14] as CCA KEM/DEM framework in this paper). Recently in [15], Kurosawa and Desmedt introduced a hybrid encryption scheme which is a modification of the hybrid scheme presented in [10]. Their scheme is interesting from both a theoretical and a practical point of view. When one looks at it as a KEM/DEM scheme, we do not know if their KEM is CCA-secure, yet the resulting scheme is CCA-secure and more efficient than the one in [10] both in computation and bandwidth. Thus the Kurosawa-Desmedt scheme points out that to obtain CCA-secure hybrid encryption, requiring both KEM/DEM to be CCA-secure, while being a sufficient condition, may not be a necessary one, and might indeed be an overkill.

1.1 Our Contributions

We present an efficient chosen cipher-text secure hybrid encryption scheme which is modification of the schemes in [15]. The way we prove the new scheme is different from [15]. We first transit the IND-CCA2 game into a IND-CPA game. Thus it does not need to answer the decryption oracle queries in the following transitions.

2 Preliminaries

We review the standard definitions of public-key encryption schemes and their security against adaptive chosen cipher-text attacks. This is followed by the definition of DDH, one-time symmetric-key encryption and one-time message authentication code.

In describing probabilistic processes, we write $x \stackrel{R}{\leftarrow} X$ to denote the action of assigning to the variable x a value sampled according to the distribution X. If S is a finite set, we simply write $s \stackrel{R}{\leftarrow} S$ to denote assignment to s of an element sampled from uniform distribution on S. If A is a probabilistic algorithm and x an input, then $A(x)$ denotes the output distribution of A on in put x. Thus, we write $y \stackrel{R}{\leftarrow} A(x)$ to denote of running algorithm A on input x and assigning the output to the variable y.

We write

$$Pr[x_1 \stackrel{R}{\leftarrow} X_1, x_2 \stackrel{R}{\leftarrow} X_2, \ldots, x_n \stackrel{R}{\leftarrow} X_n : \phi(x_1, \ldots, x_n)]$$

to denote the probability that when x_1 is drawn from a certain distribution X_1, and x_2 is drawn from a certain distribution $X_2(x_1)$, possibly depending on the particular choice of x_1, and so on, all the way to x_n, the predicate $\phi(x_1, ..., x_n)$ is true. We allow the predicate ϕ to involve the execution of probabilistic algorithms.

2.1 Public-Key Encryption

Definition 1 *A public-key encryption scheme PKE is a triple of PPT (probabilistic polynomial time) algorithms $(\mathcal{G}, \mathcal{E}, \mathcal{D})$*

- $\mathcal{G}(1^k)$: The randomized key generation algorithm $\mathcal{G}$ takes as input a security parameter (1^k) and outputs a public key PK and secret key SK. We write $(PK, SK) \leftarrow \mathcal{G}(1^k)$

- $\mathcal{E}_{PK}(m)$: The randomized encryption algorithm takes as input a public key PK and a message m, and outputs a cipher-text C. We write $C \leftarrow \mathcal{E}_{PK}(m)$
- $\mathcal{D}_{SK}(C)$: The decryption algorithm $\mathcal{D}$ takes as input a cipher-text C and secret key SK. It returns a message or the distinguished symbol $\perp$. We write $m \leftarrow \mathcal{D}_S K(C)$.

We require that for all PK, SK output by $\mathcal{G}$, all $m \in \{0,1\}^*$, and all C output by $\mathcal{E}_{PK}(m)$ we have $\mathcal{D}_{SK}(C) = m$.

We recall the standard definition of security for public-key encryption schemes against adaptive chosen cipher-text attacks.

Definition 2 *A PKE scheme is secure against adaptive chosen cipher-text attacks if the advantage of any PPT adversary A in the following game is negligible in the security parameter k:*

1. $\mathcal{G}(1^k)$ outputs PK, SK. Adversary A is given 1^k and PK.
2. The adversary may make polynomial queries to a decryption oracle $\mathcal{D}_{SK}(\cdot)$.
3. At some point, A outputs two messages m_0, m_1 with $|m_0| = |m_1|$. A bit b is randomly chosen and the adversary is given a "challenge cipher-text" $C^* \leftarrow \mathcal{E}_{PK}(m_b)$.
4. A may continue to query its decryption oracle $\mathcal{D}_{\mathcal{SK}}(\cdot)$ except that it may not request the decryption of C^*.
5. Finally, A outputs a guess b' .

   We say A succeeds if $b' = b$, and denote the probability of this event by $Pr_{A,PK}[Succ]$. The adversary's advantage is defined as $AdvCCA_A = |Pr_{A,PK}[Succ] - 1/2|$.

2.2 The Decision Diffie-Hellman Problem

There are several equivalent formulations of the decision Diffie-Hellman problem. The one that we shall use is the following.

Let G be a group of large prime order q, and consider the following two distributions:

The distribution R of random quadruples $(g_1, g_2, u_1, u_2) \in G^4$

The distribution D of quadruples $(g_1, g_2, u_1, u_2) \in G^4$, where g_1, g_2 are random, and $u_1 = g_1^r, u_2 = g_2^r$ for random $r \in Z_q$.

An algorithm that solves the decision Diffie-Hellman problem is a statistical test that can effectively distinguish these two distributions. That is, given a quadruple coming from one of the two distributions, it should output 0 or 1, and there should be a non-negligible difference between (a) the probability that it output a 1 given an input from R, and (b) the probability that it output a 1 given an input from D. The decision Diffie-Hellman problem is hard if there is no such polynomial-time statistical test.

3 Efficient hybrid encryption scheme

Our new scheme can be described as follow:

- KeyGen(1^k): Assume that G is a group of order q where q is large prime number.Constructs passive attack secure one-time symmetric-key encryption scheme SKE and forgery attack secure one-time authentication code scheme MAC as in [12]. Choose a target collision resist hash function TCR(see [12] for definition). Let $H_e : G \to \{0,1\}^{l_e}$, $H_v : G \to \{0,1\}^{l_v}$ be two hash functions, where l_e is key length of SKE, l_v is the key length of MAC. We assume that $H_e(v)$ is uniformly distributed over $\{0,1\}^{l_e}$ if v is uniformly distributed over G, $H_v(v)$ is uniformly distributed over $\{0,1\}^{l_v}$ if v is uniformly distributed over G.

$$g \overset{R}{\leftarrow} G; x, y \overset{R}{\leftarrow} Z_q; c \leftarrow g^x; d \leftarrow g^y;$$

$$PK \leftarrow (g, c, d, H_e, H_v, TCR, SKE, MAC); SK \leftarrow (x, y)$$

- Encrypt(PK, m): The encryption algorithm works as follow:

$$r \overset{R}{\leftarrow} Z_q; u \leftarrow g^r; a \leftarrow TCR(u, c^r, m); k_e \leftarrow H_e(c^r); k_v \leftarrow H_v(c^r d^{ra})$$

$$\chi \leftarrow SKE.Encrypt(k_e, m)$$

$$\tau \leftarrow MAC(k_v, \chi)$$

$$C \leftarrow (u, \chi, \tau)$$

- Decrypt(SK, C): The decryption algorithm works as follow:

$$k_e \leftarrow H_e(u^x); m \leftarrow SKE.Decrypt(k_e, \chi)$$

$$a \leftarrow TCR(u, u^x, m); k_v \leftarrow H_v(u^{x+ya})$$

$$if\ \tau \leftarrow MAC(k_v, \chi)\ return\ m; else\ return\ \perp$$

SKE is secure against passive attack, MAC is secure against forge attack. We will prove that the scheme above is CCA secure in standard model.

4 Secure proof

Remark: The proof of the security of the new scheme is different from that of Cramer and Shoup [15]. We first transit the IND-CCA2 game into a IND-CPA game. And then it does not need to answer the decryption oracle queries when transiting from Game2 to Game3.

Now we prove that our scheme is secure against chosen cipher-text attack:

Theorem 1 *The hybrid encryption scheme is secure against adaptive chosen cipher-text attack assuming that (1) the decision Diffie-Hellman problem is hard in group G, (2) SKE is secure against passive attack, and (3)MAC is secure against forgery attack, (4)TCR is secure against target collusion resist attack.*

Let game G_0 be the original attack game, let $b' \in \{0,1\}$ denote the output of A, and let T_0 be the event that $b = b'$ in G_0, so that $AdvCCA_A = |Pr[T_0] - 1/2|$.

We will define a sequence $G_1, G_2, \ldots, G_l$ of modified attack games. Each of the games operates on the same underlying probability space. In particular, the public key and secret key of the cryptosystem, the coin tosses $Coins$ of A, and the hidden bit b take on identical values across all games. For any $1 \leq i \leq l$, we let T_i be the event that $b = b'$ in game G_i.

Game G_0

1. The key generation oracle work as follow:

   I1: $g \overset{R}{\leftarrow} G; (x,y) \overset{R}{\leftarrow} Z_q$;
   I2: $c \leftarrow g^x; d \leftarrow g^y$;
   I2: $PK \leftarrow (g, c, d, H_e, H_v, TCR, SKE, MAC); SK \leftarrow (x,y)$

   Finally PK is given to A.

2. A make a sequence of calls to the decryption oracle with $C_1 \cdots C_{Q_{A1}}$. The decryption oracle computes:

   D1: $(k_{ei}) \leftarrow H_e(u_i^x)$;
   D2: $m_i \leftarrow SKE.Decrypt(k_{ei}, \chi_i)$
   D3: $a_i \leftarrow TCR(u_i, u_i^x, m_i); (k_{vi}) \leftarrow H_v(u_i^{x+ya_i})$
   D4: If $\tau_i = MAC(k_{vi}, \chi_i)$ responds m_i to A; else responds $\perp$.

3. At some point, A makes a call to the encryption oracle:

   E1: $M_0, M_1 \leftarrow A(Coins, PK, \vec{C_1}, \vec{m_1})$
   E2: $b \overset{R}{\leftarrow} \{0,1\}$
   E3: $r^* \overset{R}{\leftarrow} Z_q$
   E4: $u^* \leftarrow g^{r^*}$
   E5: $k_e^* \leftarrow H_e(c^{r^*})$;
   E6: $\chi^* \leftarrow SKE.Encrypt(k_e^*, M_b)$
   E7: $a^* \leftarrow TCR(u, c^{r^*}, M_b); k_v \leftarrow H_v(c^{r^*} d^{r^* a^*})$
   E3: $\tau^* \leftarrow MAC(k_v^*, \chi^*)$
   E9: $C^* \leftarrow (u^*, \chi^*, \tau^*)$

   where $\vec{C_1} = (C_1, \cdots, C_{Q_{A1}}), \vec{m_1} = (m_1, \cdots, m_{Q_{A1}})$

4. A may continue to query its decryption oracle with $C_{Q_{A1}+1}, \cdots, C_{Q_{A1}+Q_{A2}}$ except that it may not request the decryption of C^*. The decryption oracle works as step 2.

5. Finally, A outputs a guess b'.
   $b' \leftarrow A(Coins, PK, \vec{C}, \vec{m}, C^*, M_0, M_1)$
   where $\vec{C} = (C_1, \cdots, C_{Q_{A2}+Q_{A1}}), \vec{m} = (m_1, \cdots, m_{Q_{A2}+Q_{A1}})$

Game G_1 We remove the decryption oracle in game G_0(step 2 and step 4) to obtain a new game G_1. Let R_1 be the event that $\tau_i = MAC(c^{r_i}d^{r_i a_i})$ and A does not know m_i. This is the only way A can get help from the simulator. When A knows m_i the decryption oracle will not help him. When $\tau_i \neq MAC(c^{r_i}d^{r_i a_i})$ the simulator will response $\perp$, that's a response A knows when he try to forge a ciphertext. There are three cases we consider when R_1 happens:

Case1: In step2, $1 \leq i \leq Q_{A1}$. Since A does not know m_i, he can not calculate $a_i \leftarrow H(u_i, u_i^{z_i}, m_i)$. The value of $c^{r_i}d^{r_i a_i}$ is uniform over G and $MAC(c^{r_i}d^{r_i a_i})$ is uniform over $\{0,1\}^{l_v}$ from A's view. So the probability of this case happens is $1/q$.

Case2: In step4,$(u_i, \chi_i) \neq (u^*, \chi^*), a_i = a^*$: This case is just a target collision attack on TCR. So the probability of this case happens is $AdvTCR$. Where $AdvTCR$ is the advantage of attack on target collision resistant hash function, see [12] for detail.

Case3: In step4,$(u_i, \chi_i) \neq (u^*, \chi^*), a_i \neq a^*$, Since A does not know m_i, there are two cases we need to consider:

(a) A construct a ciphertext C_i without the help of C^*. This is the same as Case1. So the probability of this case happens is $1/q$.

(b) A construct a ciphertext C_i with the help of C^*. We have:

$$\log_g c^{r^*}d^{r^* a^*} = (x + ya^*)r^*$$

$$\log_g c^{r_i}d^{r_i a_i} = (x + ya_i)r_i$$

Since $a_i \neq a^*$, the two equations above are linear independent to each other. A can not get $c^{r_i}d^{r_i a_i}$ from $c^{r^*}d^{r^* a^*}$. If $\tau_i = MAC(c^{r_i}d^{r_i a_i})$ there is a forgery attack on MAC. So the probability of this case happens is also $AdvForge$. Where $AdvForge$ is the advantage of attack on MAC, see [12] for detail.

Now we get the probability of R_1:

$$Pr[R_1] \leq AdvTCR + AdvForge + Q_A/q$$

Where $Q_A = Q_{A1} + Q_{A2}$

Let's consider $\neg R_1$. If $\tau_i \neq c^{r_i}d^{r_i a_i}$ the decryption oracle will return $\perp$ to A. Now we see that when $\neg R_1$ happens, the decryption will not leak more information to A than he knows. And we have that:

$$Pr[T_0] = Pr[b = b' : b' \leftarrow A(Coins, PK, \vec{C}, \vec{m}, C^*, M_0, M_1)]$$

$$Pr[T_1] = Pr[b = b' : b' \leftarrow A(Coins, PK, C^*, M_0, M_1)]$$

$$Pr[T_0 \wedge \neg R_1] = Pr[b = b' : b' \leftarrow A(Coins, PK, C^*, M_0, M_1)] = Pr[T_1 \wedge \neg R_1]$$

According to [17] $lemma1$ we have:

$$|Pr[T_1] - Pr[T_0]| = Pr[R_1] \leq AdvTCR + AdvForge + Q_A/q \tag{1}$$

Game G_2 We modify the encryption oracle in game G_1 to obtain a new game G_2. Instead of using the original encryption algorithm, we modify the encryption algorithm, replacing step$E5$ with:

E5': $k_e^* \leftarrow H_e(g^{r^* x})$;

The change we have made is purely conceptual: the value of χ^* is exactly the same in game $G2$ and $G1$. Therefore,

$$Pr[T_2] = Pr[T_1] \tag{2}$$

Game G_3 We modify the encryption oracle in game G_2 to obtain a new game G_3. Instead of using the original encryption algorithm, we modify the encryption algorithm, replacing step$E5'$ with:

E5": $r^+ \stackrel{R}{\leftarrow} Z_q; k_e^* \leftarrow H_e(h^{r^+})$;

Observe that, games $G2$ and $G3$ are the same except that in game $G2$, the quadruple $(g, h, u^*, g^{r^* x})$ is uniformly distributed in D, and in game $G3$, the quadruple (g, h, u^*, g^{r^+}) is uniformly distributed in R. Thus, any difference in behavior between these two games immediately yields a statistical test for distinguishing Diffie-Hellman quadruples from random quadruples. More precisely, we have

$$|Pr[T_3] - Pr[T_2]| \leq AdvDDH \tag{3}$$

To complete the proof, we need to argue that b' is independent of b. First observe that by construction, b is independent of $Conins, PK$. In the changed encryption oracle:

$$r^* \stackrel{R}{\leftarrow} Z_q; u^* \leftarrow g^{r^*}$$

$$r^+ \stackrel{R}{\leftarrow} Z_q; k_e^* \leftarrow H_e(c^{r^+}); \chi^* \leftarrow SKE.Encrypt(k_e^*, M_b)$$

$$a^* \leftarrow TCR(u^*, c^{r^+}, M_b); k_v \leftarrow H_v(c^{r^*} d^{r^* a}); \tau^* \leftarrow MAC(k_v^*, \chi^*)$$

$$C^* \leftarrow (u^*, \chi^*, \tau^*)$$

So we get that b is independent to $C^* = (u^*, \chi^*, \tau^*)$. Now we have that b is independent to $b' \leftarrow A(Coins, PK, C^*, M_0, M_1)$. Since $Pr[T_2] = Pr[b = b' : b' \leftarrow A(Coins, PK, C^*, M_0, M_1)]$, we have:

$$|Pr[T_3] - 1/2| = 0 \tag{4}$$

From $(1), (2), (3)$ we have

$$AdvCCA_A = |Pr[T_0] - 1/2| \leq AdvDDH + AdvTCR + AdvForge + Q_A/q \tag{5}$$

Now we've proved theorem 1.

5 Efficiency Analysis

Shoup's [10] hybrid public-key encryption scheme is the first hybrid encryption scheme in standard model. Kurosawa and Desmedt's [15] (KD) hybrid encryption scheme is a modification of Shoup's scheme. The efficiency of our schemes , Shoup and KD is list in table 1.

Table 1: Efficiency comparison for CCA-secure public-key encryption schemes Scheme

	Encryption(exp)	Decryption(exp)	Cipher-text overhead(bit)
Shoup	4.5(3exp+1mexp)	2.5(1exp+1mexp)	$3\|q\| + \|\tau\|$
KD	3.5(2exp+1mexp)	1.5(1mexp)	$2\|q\| + \|\tau\|$
NEW	3(3exp)	2(2exp)	$\|q\| + \|\tau\|$

Where NEW is our hybrid encryption scheme above. When tabulating computational efficiency hash function and block cipher evaluations are ignored, multi-exponentiation ($mexp$) is counted as 1.5 exponentiations (exp). Cipher-text overhead represents the difference between the cipher-text length and the message length, and $|q|$ is the length of a group element, $|\tau|$ is the length of tag in MAC.

6 Conclusion

Kurosawa-Desmedt's scheme shows out that to obtain CCA-secure hybrid encryption, requiring both KEM/DEM to be CCA-secure, while being a sufficient condition, may not be a necessary one, and might indeed be an overkill. By modified the scheme of KD we get a more efficient hybrid encryption scheme,which is provably secure in standard model under DDH assumption. The new scheme is more efficient than Kurosawa-Desmedt's scheme on encryption and ciphertext overhead.

References

[1] M. Bellare, A. Desai, D. Pointcheval, and P. Rogaway, "Relations Among Notions of Security for Public-Key Encryption Schemes", *Adv. in Cryptology - Crypto 1998*, LNCS vol. 1462, Springer-Verlag, pp. 26-45, 1998;

[2] D. Dolev, C. Dwork, and M. Naor, "Non-Malleable Cryptography", *SIAM J. Computing* , 30(2): 391-437, 2000;

[3] C. Rackoff and D. Simon, "Non-Interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack", *Adv. in Cryptology - Crypto 1991*, LNCS vol. 576, Springer- Vrlag , pp. 433-444, 1991;

[4] V. Shoup, "Why Chosen Ciphertext Security Matters", *IBM Research Report RZ 3076*, November , 1998. Available at http://www.shoup.net/papers;

[5] V. Shoup, "A Proposal for an ISO Standard for Public Key Encryption (version 2.1)", December, 2001. Available at http://www.shoup.net/papers;

[6] R. Canetti, O. Goldreich, and S. Halevi, "The Random Oracle Methodology Revisited" *30thACM Symp. on Theory of Computing (STOC), ACM*, pp. 209-218, 1998;

[7] R. Cramer and V. Shoup, "A Practical Public Key Cryptosystem Provably Secure Against Chosen Ciphertext Attack", *Adv. in Cryptology - Crypto 1998*, LNCS vol. 1462, Springer-Verlag , pp. 13-25, 1998;

[8] R. Cramer and V. Shoup, "Universal Hash Proofs and a Paradigm for Adaptive Chosen Ciphertext Secure Public-Key Encryption", *Adv. in Cryptology - Eurocrypt 2002*, LNCS vol. 2332,Springer-Verlag, pp. 45-64, 2002;

[9] M. Abdalla, M. Bellare and P. Rogaway. DHIES: An encryption scheme based on the Diffie-Hellman Problem Extended abstract, entitled The Oracle Diffie-Hellman Assumptions and an Analysis of DHIES, was in Topics in Cryptology - CT-RSA 01, Lecture Notes in Computer Science Vol. 2020, D. Naccache ed, Springer-Verlag, 2001

[10] V. Shoup. Using hash functions as a hedge against chosen ciphertext attack. In B. Preneel, editor, Advances in Cryptology - Eurocrypt 2000, volume 1807 of Lecture Notes in Computer Science, pages 275-288. Springer-Verlag, 2000.

[11] T. Okamoto and D. Pointcheval. REACT: Rapid enhanced-security asymmetric cryptosystem transform. In RSA 2001, LNCS, Springer-Verlag, 2001.

[12] R. Cramer and V. Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM Journal on Computing, 33(1):167-226, 2003.

[13] International Organization for Standardization. ISO/IEC CD 18033- 2, Information technology – Security techniques – Encryption Algo- rithms – Part 2: Asymmetric Ciphers, 2003.

[14] V. Shoup. ISO 18033-2: An emerging standard for public-key encryption (committee draft). Available at http://shoup.net/iso/, June 3 2004.

[15] K. Kurosawa and Y. Desmedt. A new paradigm of hybrid encryp- tion scheme. In M. Franklin, editor, Advances in Cryptology - Crypto 2004, volume 3152 of Lecture Notes in Computer Sciene, pages 426-442. Springer-Verlag, 2004.

[16] Masayuki Abe, Rosario Gennaro, Kaoru Kurosawa, and Victor Shoup, Tag-KEM/DEM: A new framework for hybrid encryption and a new analysis of Kurosawa-Desmedt KEM", by Abe, Gennaro, Kurosawa, and Shoup, in Proc. Eurocrypt 2005.

[17] Victor Shoup, Sequences of games: a tool for taming complexity in security proofs, manuscript, Nov. 30, 2004. Revised, May 27, 2005; Jan. 18, 2006. http://shoup.net/papers/