

An Adaptation of the NICE Cryptosystem to Real Quadratic Orders

Michael J. Jacobson, Jr.^{1,*}, Renate Scheidler², and Daniel Weimer³

¹ Department of Computer Science, University of Calgary,
2500 University Drive NW, Calgary, Alberta T2N 1N4, Canada
`jacobsc@cpsc.ucalgary.ca`

² Department of Mathematics & Statistics, University of Calgary,
2500 University Drive NW, Calgary, Alberta T2N 1N4, Canada
`rscheidl@math.ucalgary.ca`

³ Charles River Development, 7 New England Executive Park,
Burlington MA 01803 USA
`daniel.weimer@gmx.de`

Abstract. In 2000, Paulus and Takagi introduced a public key cryptosystem called NICE that exploits the relationship between maximal and non-maximal orders in imaginary quadratic number fields. Relying on the intractability of integer factorization, NICE provides a similar level of security as RSA, but has faster decryption. This paper presents REAL-NICE, an adaptation of NICE to orders in real quadratic fields. REAL-NICE supports smaller public keys than NICE, and while preliminary computations suggest that it is somewhat slower than NICE, it still significantly outperforms RSA in decryption.

1 Introduction

The most well-known and widely used public-key cryptosystem whose security is related to the intractability of the integer factorization problem is the RSA scheme. A lesser known factoring-based system is the NICE (**N**ew **I**deal **C**oset **E**ncryption) scheme [13,18], a cryptosystem whose trapdoor decryption makes use of the relationship between ideals in the maximal and a non-maximal order of an imaginary quadratic number field. The security of NICE relies on the presumed intractability of factoring an integer of the form q^2p where p and q are prime, thereby providing a similar level of security as RSA, but with much faster decryption. NICE decryption has quadratic complexity, as opposed to RSA's cubic decryption complexity. This makes NICE particularly suited for devices with limited computing power or applications that require fast digital signature generation.

In this paper, we explain how to extend the NICE concept to real quadratic fields; this was first proposed in [19]. REAL-NICE exploits the same relationship between ideals in the maximal and a non-maximal quadratic order as NICE. Furthermore, just as in NICE, knowledge of the trapdoor information is provably

* Research by the first two authors supported by NSERC of Canada.

equivalent to being able to factor the discriminant of the non-maximal order in random polynomial time. However, the security of REAL-NICE relies on the intractability of a somewhat different problem. In NICE, encryption hides the message ideal in its own exponentially large coset with respect to a certain subgroup of the ideal class group of the non-maximal order. In a real quadratic field, such a coset may be too small to prevent an exhaustive search attack. Instead, REAL-NICE encryption hides the message ideal in the generally exponentially large cycle of reduced ideals in its own ideal class in the non-maximal order.

While preliminary numerical data using prototype implementations suggest that REAL-NICE is somewhat slower than its imaginary counterpart NICE, REAL-NICE allows for the possibility of a smaller public key than NICE, at the expense of increased encryption effort. Moreover, both our NICE and REAL-NICE prototypes significantly outperformed a highly optimized public-domain implementation of RSA in decryption for all five NIST security levels [12]; for the two highest such levels, combined encryption and decryption was faster for both NICE and REAL-NICE compared to RSA.

The discrepancy in performance between NICE and REAL-NICE can be offset by using a more efficient encryption algorithm, called IMS encryption, for REAL-NICE. IMS encryption exploits the very fast baby step operation in the cycle of reduced ideals of a real quadratic order, an operation that has no imaginary analogue. Unfortunately, so far, the only known rigorous proof of security for IMS encryption needs to assume a very unfavourable parameter set-up. However, even under these adverse assumptions, IMS-REAL-NICE outperformed the original REAL-NICE system. It is conceivable that a set-up could be established that makes IMS-REAL-NICE competitive to NICE without sacrificing security. IMS encryption and its security are the subject of future research.

2 Overview of Quadratic Orders

We begin with a brief overview of quadratic fields and their orders. Most of the material in this section can be found in [11] and Chapter 2, §7, of [4]; while the latter source considers mostly imaginary quadratic fields, much of the results are easily extendable to real quadratic fields as was done in [19].

Let $D \in \mathbb{Z}$, $D \neq 0, \pm 1$ be a squarefree integer. A *quadratic (number) field* is a field of the form $\mathcal{K} = \mathbb{Q}(\sqrt{D}) = \{a + b\sqrt{D} \mid a, b \in \mathbb{Q}\}$. \mathcal{K} is an *imaginary*, respectively, *real* quadratic field if $D < 0$, respectively, $D > 0$. Set $\Delta_1 = 4D$ if $D \equiv 2$ or $3 \pmod{4}$ and $\Delta_1 = D$ if $D \equiv 1 \pmod{4}$, so $\Delta_1 \equiv 0$ or $1 \pmod{4}$. Δ_1 is called a *fundamental discriminant*. For $f \in \mathbb{N}$, set $\Delta_f = f^2 \Delta_1$. The (*quadratic*) *order of conductor f* in \mathcal{K} is the \mathbb{Z} -submodule \mathcal{O}_{Δ_f} of \mathcal{K} of rank 2 generated by 1 and $f(\Delta_1 + \sqrt{\Delta_1})/2$; its *discriminant* is Δ_f . We speak of imaginary, respectively, real quadratic orders, depending on whether \mathcal{K} is an imaginary, respectively, a real quadratic field. The *maximal order* of \mathcal{K} is \mathcal{O}_{Δ_1} ; it contains all the orders of \mathcal{K} , and $f = [\mathcal{O}_{\Delta_1} : \mathcal{O}_{\Delta_f}]$ is the index of \mathcal{O}_{Δ_f} in \mathcal{O}_{Δ_1} as an additive subgroup.

Henceforth, let $f \in \mathbb{N}$ be any conductor. We denote by $\mathcal{O}_{\Delta_f}^*$ the group of units of the integral domain \mathcal{O}_{Δ_f} , i.e. the group of divisors of 1 in \mathcal{O}_{Δ_f} . The units

of \mathcal{O}_{Δ_f} , denoted by $\mathcal{O}_{\Delta_f}^*$, form an Abelian group under multiplication. If \mathcal{K} is imaginary, then $\mathcal{O}_{\Delta_f}^*$ consists of the roots of unity in \mathcal{K} and thus has 6, 4, or 2 elements, according to whether $\Delta_1 = -3$, $\Delta_1 = -4$, or $\Delta_1 < -4$. If \mathcal{K} is real, then $\mathcal{O}_{\Delta_f}^*$ is an infinite cyclic group with finite torsion $\{1, -1\}$, whose unique generator ϵ_{Δ_f} exceeding 1 is the *fundamental unit* of \mathcal{O}_{Δ_f} . In this case, the real number $R_{\Delta_f} = \log(\epsilon_{\Delta_f})$ is the *regulator* of \mathcal{O}_{Δ_f} . Here, as usual, $\log(x)$ denotes the natural logarithm of $x > 0$.

An (*integral*) \mathcal{O}_{Δ_f} -ideal¹ \mathfrak{a} is a \mathbb{Z} -submodule of \mathcal{O}_{Δ_f} of rank 2 that is closed under multiplication by elements in \mathcal{O}_{Δ_f} . A *fractional* \mathcal{O}_{Δ_f} -ideal \mathfrak{a} is a \mathbb{Z} -submodule of \mathcal{K} of rank 2 such that $d\mathfrak{a}$ is an (integral) \mathcal{O}_{Δ_f} -ideal for some $d \in \mathbb{N}$. A fractional \mathcal{O}_{Δ_f} -ideal \mathfrak{a} is *invertible* if there exists a fractional \mathcal{O}_{Δ_f} -ideal \mathfrak{b} such that $\mathfrak{a}\mathfrak{b} = \mathcal{O}_{\Delta_f}$, where the product of two fractional \mathcal{O}_{Δ_f} -ideals $\mathfrak{a}, \mathfrak{b}$ is defined to consist of all finite sums of products of the form $\alpha\beta$ with $\alpha \in \mathfrak{a}$ and $\beta \in \mathfrak{b}$. The set of invertible fractional \mathcal{O}_{Δ_f} -ideals, denoted by $\mathcal{I}(\mathcal{O}_{\Delta_f})$, is an infinite Abelian group under multiplication with identity \mathcal{O}_{Δ_f} . A *principal* fractional \mathcal{O}_{Δ_f} -ideal \mathfrak{a} consists of \mathcal{O}_{Δ_f} -multiples of some fixed element $\alpha \in \mathcal{K}^* = \mathcal{K} \setminus \{0\}$ that is said to *generate* (or be a *generator* of) \mathfrak{a} . We write $\mathfrak{a} = (\alpha) = \alpha\mathcal{O}_{\Delta_f}$. The principal fractional \mathcal{O}_{Δ_f} -ideals form an infinite subgroup of $\mathcal{I}(\mathcal{O}_{\Delta_f})$ that is denoted by $\mathcal{P}(\mathcal{O}_{\Delta_f})$. The factor group $Cl(\mathcal{O}_{\Delta_f}) = \mathcal{I}(\mathcal{O}_{\Delta_f})/\mathcal{P}(\mathcal{O}_{\Delta_f})$ is a finite Abelian group under multiplication, called the *ideal class group* of \mathcal{O}_{Δ_f} . Its order h_{Δ_f} is the (*ideal*) *class number* of \mathcal{O}_{Δ_f} . For any \mathcal{O}_{Δ_f} -ideal \mathfrak{a} , we denote the \mathcal{O}_{Δ_f} -ideal class by $[\mathfrak{a}] \in Cl(\mathcal{O}_{\Delta_f})$.

For any element $\alpha = a + b\sqrt{D} \in \mathcal{K}$ ($a, b \in \mathbb{Q}$), the *conjugate* of α is $\bar{\alpha} = a - b\sqrt{D} \in \mathcal{K}$, and the *norm* of α is $N(\alpha) = \alpha\bar{\alpha} = a^2 - b^2D \in \mathbb{Q}$. If $\alpha \in \mathcal{O}_{\Delta_1}$, then $N(\alpha) \in \mathbb{Z}$. The *norm* $N_{\Delta_f}(\mathfrak{a})$ of an (integral) \mathcal{O}_{Δ_f} -ideal \mathfrak{a} is the index of \mathfrak{a} as an additive subgroup of \mathcal{O}_{Δ_f} . When the context is clear, we will omit the subscript Δ_f from the ideal norm and simply write $N(\mathfrak{a})$. If we set $\bar{\mathfrak{a}} = \{\bar{\alpha} \mid \alpha \in \mathfrak{a}\}$, then $\mathfrak{a}\bar{\mathfrak{a}} = (N(\mathfrak{a}))$, the principal \mathcal{O}_{Δ_f} -ideal generated by $N(\mathfrak{a})$. If \mathfrak{a} is a principal \mathcal{O}_{Δ_f} -ideal generated by $\alpha \in \mathcal{O}_{\Delta_f}$, then $N(\mathfrak{a}) = |N(\alpha)|$.

An integral \mathcal{O}_{Δ_f} -ideal \mathfrak{a} is *primitive* if the only positive integer d such that every element of \mathfrak{a} is an \mathcal{O}_{Δ_f} -multiple of d is $d = 1$. An \mathcal{O}_{Δ_f} -ideal \mathfrak{a} is *reduced* if it is primitive and there does not exist any non-zero $\alpha \in \mathfrak{a}$ with $|\alpha| < N(\mathfrak{a})$ and $|\bar{\alpha}| < N(\mathfrak{a})$. We summarize some important properties of reduced ideals; see for example [6,13,21] as well as Sections 2.1 and 2.2 of [19].

Theorem 1. *Let \mathcal{O}_{Δ_f} be an order in a quadratic number field $\mathcal{K} = \mathbb{Q}(\sqrt{D})$. Then the following hold:*

1. *Every ideal class of $Cl(\mathcal{O}_{\Delta_f})$ contains a reduced \mathcal{O}_{Δ_f} -ideal.*
2. *If \mathcal{K} is imaginary, then every ideal class of $Cl(\mathcal{O}_{\Delta_f})$ contains a unique reduced \mathcal{O}_{Δ_f} -ideal. If \mathcal{K} is real, then the number $r_{\mathbf{C}}$ of reduced ideals in any ideal class $\mathbf{C} \in Cl(\mathcal{O}_{\Delta_f})$ satisfies $R_{\Delta_f}/\log(f^2D) \leq r_{\mathbf{C}} < 2R_{\Delta_f}/\log(2) + 1$.*
3. *If \mathfrak{a} is a primitive \mathcal{O}_{Δ_f} -ideal with $N(\mathfrak{a}) < \sqrt{|\Delta_f|}/2$, then \mathfrak{a} is reduced.*
4. *If \mathfrak{a} is a reduced \mathcal{O}_{Δ_f} -ideal, then $N(\mathfrak{a}) < \sqrt{\Delta_f}$ if \mathcal{K} is real and $N(\mathfrak{a}) < \sqrt{|\Delta_f|/3}$ if \mathcal{K} is imaginary.*

¹ We always assume that integral and fractional ideals are non-zero.

For an \mathcal{O}_{Δ_f} -ideal \mathfrak{a} , we denote by $\rho_{\Delta_f}(\mathfrak{a})$ any reduced \mathcal{O}_{Δ_f} -ideal in the class of \mathfrak{a} . By the above theorem, if \mathcal{K} is imaginary, then $\rho_{\Delta_f}(\mathfrak{a})$ is the unique reduced representative in the \mathcal{O}_{Δ_f} -ideal class of \mathfrak{a} , whereas if \mathcal{K} is real, then there are many choices for $\rho_{\Delta_f}(\mathfrak{a})$. Given any \mathcal{O}_{Δ_f} -ideal \mathfrak{a} , a reduced ideal $\rho_{\Delta_f}(\mathfrak{a})$ in the equivalence class of \mathfrak{a} can be found using at most $O(\log(N(\mathfrak{a})/\sqrt{\Delta_f}) \log(\Delta_f))$ bit operations. Furthermore, in the real scenario, the entire cycle of reduced ideals in the \mathcal{O}_{Δ_f} -ideal class of \mathfrak{a} can then be traversed using a procedure called *baby steps*. Details on ideal reduction, baby steps, and other ideal arithmetic can be found in Section 7.

For any integer d , an integral \mathcal{O}_{Δ_f} -ideal \mathfrak{a} is said to be *prime to d* if $N(\mathfrak{a})$ is relatively prime to d . Of particular interest is the case $d = f$, as every \mathcal{O}_{Δ_f} -ideal prime to f is invertible, and the norm map is multiplicative on the set of \mathcal{O}_{Δ_f} -ideals prime to f . Denote by $\mathcal{I}(\mathcal{O}_{\Delta_f}, f)$ the subgroup of $\mathcal{I}(\mathcal{O}_{\Delta_f})$ generated by the \mathcal{O}_{Δ_f} -ideals prime to f , by $\mathcal{P}(\mathcal{O}_{\Delta_f}, f)$ the subgroup of $\mathcal{I}(\mathcal{O}_{\Delta_f}, f)$ generated by the principal ideals (α) with $\alpha \in \mathcal{O}_{\Delta_f}$ and $N(\alpha)$ prime to f , and set $Cl(\mathcal{O}_{\Delta_f}, f) = \mathcal{I}(\mathcal{O}_{\Delta_f}, f)/\mathcal{P}(\mathcal{O}_{\Delta_f}, f)$. Then $Cl(\mathcal{O}_{\Delta_f}, f)$ is isomorphic to the class group $Cl(\mathcal{O}_{\Delta_f})$ of \mathcal{O}_{Δ_f} ; see Proposition 7.19, p. 143, of [4] and Theorem 2.16, p. 10, of [19].

Finally, we denote by $\mathcal{I}(\mathcal{O}_{\Delta_1}, f)$ the subgroup of $\mathcal{I}(\mathcal{O}_{\Delta_1})$ generated by the \mathcal{O}_{Δ_1} -ideals prime to f , by $\mathcal{P}(\mathcal{O}_{\Delta_1}, f)$ the subgroup of $\mathcal{I}(\mathcal{O}_{\Delta_1}, f)$ generated by the principal \mathcal{O}_{Δ_1} -ideals (α) with $\alpha \in \mathcal{O}_{\Delta_1}$ and $N(\alpha)$ prime to f , and define the factor group $Cl(\mathcal{O}_{\Delta_1}, f) = \mathcal{I}(\mathcal{O}_{\Delta_1}, f)/\mathcal{P}(\mathcal{O}_{\Delta_1}, f)$.

For the NICE cryptosystem in both real and imaginary quadratic orders, it will be important to move between \mathcal{O}_{Δ_f} -ideals prime to f and \mathcal{O}_{Δ_1} -ideals prime to f . More specifically, we have the following isomorphism (see Proposition 7.20, p. 144, of [4] and Theorem 3.2, p. 25, of [19]):

$$\phi : \mathcal{I}(\mathcal{O}_{\Delta_1}, f) \longrightarrow \mathcal{I}(\mathcal{O}_{\Delta_f}, f) \quad \text{via} \quad \phi(\mathfrak{A}) = \mathfrak{A} \cap \mathcal{O}_{\Delta_f}, \quad \phi^{-1}(\mathfrak{a}) = \mathfrak{a} \mathcal{O}_{\Delta_1} \quad . \quad (2.1)$$

The maps ϕ and ϕ^{-1} are efficiently computable if f and Δ_1 are known; for details, see Section 7. In fact, both the NICE and the REAL-NICE schemes use ϕ^{-1} as their underlying trapdoor one-way function, with public information Δ_f and trapdoor information f , where f is a prime. Note that ϕ and ϕ^{-1} preserve norms and primitivity. Furthermore, ϕ^{-1} preserves ideal principality, but ϕ does not. Thus, ϕ^{-1} induces a surjective homomorphism

$$\hat{\Phi} : Cl(\mathcal{O}_{\Delta_f}, f) \longrightarrow Cl(\mathcal{O}_{\Delta_1}, f) \quad \text{via} \quad \hat{\Phi}([\mathfrak{a}]) = [\phi^{-1}(\mathfrak{a})] = [\mathfrak{a} \mathcal{O}_{\Delta_1}] \quad . \quad (2.2)$$

For proofs of these results, see pp. 144-146 of [4] and pp. 25-29 of [19].

The kernel of $\hat{\Phi}$, i.e. the subgroup of $Cl(\mathcal{O}_{\Delta_f}, f)$ of the form

$$\ker(\hat{\Phi}) = \{[\mathfrak{a}] \in Cl(\mathcal{O}_{\Delta_f}, f) \mid \phi^{-1}(\mathfrak{a}) \text{ is a principal } \mathcal{O}_{\Delta_1}\text{-ideal}\}$$

is of crucial importance to the NICE cryptosystem in imaginary quadratic orders, and also plays a role in its counterpart REAL-NICE in real quadratic orders. The size of this kernel is exactly the class number ratio $h_{\Delta_f}/h_{\Delta_1}$. For the cryptographically interesting case of prime conductor $f = q$, and disregarding the

small cases $\Delta_1 = -3$ or -4 where \mathcal{K} contains nontrivial roots of unity, the size of this kernel is given by

$$|\ker(\hat{\Phi})| = \frac{h_{\Delta_q}}{h_{\Delta_1}} = \begin{cases} q - (\Delta_1/q) & \text{if } \Delta_1 < -4, \\ q - (\Delta_1/q)R_{\Delta_1}/R_{\Delta_q} & \text{if } \Delta_1 > 0, \end{cases} \quad (2.3)$$

where (Δ_1/q) denotes the Legendre symbol.

3 The Original NICE Cryptosystem

The original NICE cryptosystem [18,13] exploits the relationship between ideals in a maximal and a non-maximal imaginary quadratic order of prime conductor q as described in (2.1) and (2.2). The key observation is that images of \mathcal{O}_{Δ_q} -ideals under the map ϕ^{-1} of (2.1) are efficiently computable if q is known, whereas without knowledge of the trapdoor information q (i.e. only knowledge of Δ_q), this task is infeasible and is in fact provably equivalent to being able to factor Δ_q in random polynomial time (see Theorem 2.1, pp. 13-14, of [18]).

The specifics of NICE are as follows:

Private Key: Two large primes p, q of approximately equal size with $p \equiv 3 \pmod{4}$.

Public Key: $(\Delta_q, k, n, \mathfrak{p})$ where

- $\Delta_q = q^2 \Delta_1$ with $\Delta_1 = -p$;
- k and n are the bit lengths of $\lfloor \sqrt{|\Delta_1|}/4 \rfloor$ and $q - (\Delta_1/q)$, respectively;
- \mathfrak{p} is a randomly chosen \mathcal{O}_{Δ_q} -ideal with $[\mathfrak{p}] \in \ker(\hat{\Phi})$.

The key ideal \mathfrak{p} can be found by generating a random element $\alpha \in \mathcal{O}_{\Delta_1}$ whose norm is not divisible by q , finding a \mathbb{Z} -basis of the principal \mathcal{O}_{Δ_1} -ideal $\mathfrak{A} = (\alpha)$, and computing $\mathfrak{p} = \phi(\mathfrak{A})$. Note that the \mathcal{O}_{Δ_q} -ideal \mathfrak{p} itself is generally not principal, but its image $\phi^{-1}(\mathfrak{p})$ is a principal \mathcal{O}_{Δ_1} -ideal.

Encryption: Messages are bit strings of bit length $k - t$, where t is a fixed parameter explained below. To encrypt a message m :

1. Embed m into a primitive \mathcal{O}_{Δ_q} -ideal \mathfrak{m} prime to q with $N_{\Delta_q}(\mathfrak{m}) \leq 2^k$ in such a way that $N_{\Delta_q}(\mathfrak{m})$ uniquely determines m .
2. Generate random $r \in_R \{1, 2, \dots, 2^{n-1}\}$.
3. The ciphertext is the reduced \mathcal{O}_{Δ_q} -ideal $\mathfrak{c} = \rho_{\Delta_q}(\mathfrak{m}\mathfrak{p}^r)$.

Note that since $2^{n-1} < q - (\Delta_1/q) < 2^n$, the range for r specified in step 2 ensures that $r < q - (\Delta_1/q) = |\ker(\hat{\Phi})|$. This is the optimal range, as $[\mathfrak{p}]^{q - (\Delta_1/q)}$ is the identity in $Cl(\mathcal{O}_{\Delta_q})$, i.e. the principal class. The cipher ideal \mathfrak{c} is computed using standard ideal arithmetic; see Section 7 for details.

Decryption: To decrypt a ciphertext \mathcal{O}_{Δ_q} -ideal \mathfrak{c} :

1. Compute $\mathfrak{M} = \rho_{\Delta_1}(\phi^{-1}(\mathfrak{c}))$.
2. Extract m from $N_{\Delta_1}(\mathfrak{M})$.

Note that $[\mathfrak{p}] \in \ker(\hat{\Phi})$ and (2.2) together imply

$$\begin{aligned} [\mathfrak{M}] &= [\rho_{\Delta_1}(\phi^{-1}(\mathfrak{c}))] = [\phi^{-1}(\mathfrak{c})] = \hat{\Phi}([\mathfrak{c}]) = \hat{\Phi}([\rho_{\Delta_q}(\mathfrak{m}\mathfrak{p}^r)]) \\ &= \hat{\Phi}([\mathfrak{m}\mathfrak{p}^r]) = \hat{\Phi}([\mathfrak{m}])\hat{\Phi}([\mathfrak{p}])^r = \hat{\Phi}([\mathfrak{m}]) = [\phi^{-1}(\mathfrak{m})] . \end{aligned}$$

Since ϕ^{-1} is norm-preserving and $2^{k-1} \leq \lfloor \sqrt{|\Delta_1|}/4 \rfloor < 2^k$, encryption step 1 yields

$$N_{\Delta_1}(\phi^{-1}(\mathfrak{m})) = N_{\Delta_q}(\mathfrak{m}) \leq 2^k \leq 2 \left\lfloor \frac{\sqrt{|\Delta_1|}}{4} \right\rfloor < \frac{\sqrt{|\Delta_1|}}{2} ,$$

where the last inequality follows since $\sqrt{|\Delta_1|}/4 \notin \mathbb{Z}$. By part 3 of Theorem 1, $\phi^{-1}(\mathfrak{m})$ is a reduced \mathcal{O}_{Δ_1} -ideal. Thus, \mathfrak{M} and $\phi^{-1}(\mathfrak{m})$ are reduced ideals in the same \mathcal{O}_{Δ_1} -ideal class, so they must be equal by part 2 of Theorem 1. It follows that $N_{\Delta_1}(\mathfrak{M}) = N_{\Delta_1}(\phi^{-1}(\mathfrak{m})) = N_{\Delta_q}(\mathfrak{m})$, which by encryption step 1 uniquely determines m . Note also that $N_{\Delta_f}(\mathfrak{m}) = N_{\Delta_1}(\mathfrak{M}) < \sqrt{|\Delta_1|}/2 < q$, where the last inequality holds because p and q are of roughly the same size. It follows that both \mathfrak{m} and \mathfrak{M} are prime to q . Since $N(\mathfrak{c}) = N(\mathfrak{m})N(\mathfrak{p})^r$, \mathfrak{c} is also prime to q .

Since the decrypter knows the conductor q of \mathcal{O}_{Δ_q} , he can efficiently compute $\phi^{-1}(\mathfrak{c})$, and hence \mathfrak{M} using standard reduction arithmetic. We explain how to compute images under ϕ^{-1} in Section 7.

To perform encryption step 1, one first selects a security parameter t ; we explain below how large t should be chosen. The plaintext needs to be divided into message blocks of bit length $k - t$. To embed such a block m into a reduced \mathcal{O}_{Δ_q} -ideal \mathfrak{m} prime to q , the encrypter does the following:

1. Set $\overline{m} = m2^t$, obtaining an integer \overline{m} of bit length k whose t low order bits are all 0.
2. Find the smallest prime l exceeding \overline{m} such that $(\Delta_q/l) = 1$.
3. Set \mathfrak{m} to be the \mathcal{O}_{Δ_q} -ideal of norm l .

If $l \equiv 3 \pmod{4}$, then a \mathbb{Z} -basis for the ideal \mathfrak{m} can be found efficiently and deterministically. If $l \equiv 1 \pmod{4}$, then there is a fast probabilistic method for performing step 3 above. For details, see again Section 7.

If $l \leq \overline{m} + 2^t$, then $m \leq 2^{k-t} - 1$ implies $N_{\Delta_f}(\mathfrak{m}) = l \leq \overline{m} + 2^t = (m+1)2^t \leq 2^k$ as desired. Furthermore, the k high order bits of l agree with \overline{m} and hence with m . Since $N_{\Delta_1}(\mathfrak{M}) = N_{\Delta_q}(\mathfrak{m}) = l$, m is easily obtained from $N_{\Delta_1}(\mathfrak{M})$ in decryption step 2 by truncating the first k bits from l . According to pp. 34-36 of [19], the probability that $l \leq \overline{m} + 2^t$ is bounded below by $P_t = 1 - 2^{-2^t/k}$. It follows that decryption step 2 is successful with high probability for t sufficiently large.

The security of NICE was analyzed in detail in [13], [7], and [19], and resides in the difficulty of factoring Δ_q . We only briefly review some facts here. Encryption under NICE can be viewed as masking the message ideal \mathfrak{m} by multiplying it by a random ideal $\mathfrak{a} = \mathfrak{p}^r$ with $[\mathfrak{a}] \in \ker(\hat{\Phi})$, thereby hiding it in its own coset $\mathfrak{m} \ker(\hat{\Phi})$. The size of each such coset is equal to $|\ker(\hat{\Phi})| = q - (\Delta_1/q)$. Obviously, q must be chosen large enough to make exhaustive search through any coset relative to $\ker(\hat{\Phi})$ infeasible. Moreover, in order to guarantee a sufficiently large number of

distinct elements of the form $\mathfrak{m}\mathfrak{p}^r$, or equivalently, a sufficiently large number of distinct powers \mathfrak{p}^r , we need to ensure that the subgroup in $\ker(\hat{\Phi})$ generated by the class $[\mathfrak{p}]$ is large. The order of this subgroup is a divisor of $q - (\Delta_1/q)$, so this quantity should be chosen prime or almost prime. Suppose that $q - (\Delta_1/q) = Ld$ where L is a large prime and $d \in \mathbb{N}$ is very small. Then the number of generators of the cyclic subgroup of $\ker(\hat{\Phi})$ of order L is $\phi(L) = L - 1$, where $\phi(N)$ denotes the Euler totient function of $N \in \mathbb{N}$. So the probability that a random ideal $\mathfrak{p} \in \ker(\hat{\Phi})$ generates this subgroup is $(L - 1)/Ld \approx 1/d$ which is large if d is small. One expects d trials of an ideal \mathfrak{p} to produce a desirable key ideal. For any such trial, checking that $\rho_{\Delta_q}(\mathfrak{p}^d) \neq \mathcal{O}_{\Delta_q}$ guarantees that \mathfrak{p} generates a subgroup of $\ker(\hat{\Phi})$ of order L .

An algorithm for computing images of primitive \mathcal{O}_{Δ_q} -ideals under ϕ^{-1} without knowledge of q would lead to the decryption of any message. However, according to Theorem 1 of [13], such an algorithm could be used as an oracle for factoring Δ_q in random polynomial time. Hence, the security of NICE is equivalent to factoring an integer of the form q^2p , so p and q need to be chosen sufficiently large to render the factorization of Δ_q via the elliptic curve method and the number field sieve infeasible. Using the estimate that factoring a 1024-bit RSA modulus is computationally equivalent to finding a 341-bit factor of a 3-prime modulus of the same size [10] yields the estimates in Table 1 for parameter sizes of Δ_q that are required to provide a level of security equivalent to block ciphers with keys of 80, 112, 128, 192, and 256 bits, respectively.

Table 1. NIST recommendations for parameter sizes of p and q

symmetric key size	80	112	128	192	256
Size of Δ_q	1024	2048	3072	8192	15360
Size of p and q	341	682	1024	2731	5120

To the best of our knowledge, revealing the public key and the form of Δ_q ($\Delta_q = -q^2p$ with primes p, q and $p \equiv 3 \pmod{4}$) does not compromise the security of NICE. Finally, the chosen ciphertext attack of [7] is prevented by padding message blocks with t low order 0 bits for sufficiently large t . To ensure that this attack is approximately as costly as any other known attack method, t should be chosen according to the first row of Table 1 (the symmetric key size). NICE has also been extended, using standard techniques, to provide IND-CCA2 security in the random oracle model — see the NICE-X protocol presented in [1].

4 NICE in Real Quadratic Orders

Before we describe in detail REAL-NICE, our adaptation of NICE to orders in real quadratic fields, we highlight the main differences between REAL-NICE and NICE. The security of the original NICE scheme resides in the difficulty of

identifying a specific representative in the coset of an \mathcal{O}_{Δ_q} -ideal \mathfrak{m} relative to $\ker(\hat{\Phi})$ without knowledge of the conductor q of the order \mathcal{O}_{Δ_q} . In real quadratic orders, this problem is generally easy to solve via exhaustive search, since h_{Δ_q} can be a small multiple of h_{Δ_1} , resulting in a very small kernel of $\hat{\Phi}$ by (2.3). Instead, in the real case, an adversary needs to identify a specific reduced ideal in the cycle of reduced ideals in the \mathcal{O}_{Δ_q} -ideal class $[\mathfrak{m}]$. It is therefore necessary to ensure that the number of reduced ideals in any \mathcal{O}_{Δ_q} -ideal class is large.

At the same time, the decryption process of NICE no longer yields a unique reduced \mathcal{O}_{Δ_1} -ideal. To extract m , we need to make sure that the \mathcal{O}_{Δ_1} -ideal class of $\phi^{-1}(\mathfrak{c})$ contains very few reduced ideals, so they can all be quickly computed and the correct one identified by a predetermined unique bit pattern in its norm. During encryption, m is endowed with that same bit pattern. By part 2 of Theorem 1, the system parameters must therefore be chosen so that R_{Δ_1} is very small, while R_{Δ_q} is large.

Finally, the ideal \mathfrak{p} need no longer be included in the public key; instead, a random ideal \mathfrak{p} with $[\mathfrak{p}] \in \ker(\hat{\Phi})$ can be generated for each encryption.

The specifics of REAL-NICE are as follows:

Private Key: Two large primes p, q of approximately equal size with $p \equiv 1 \pmod{4}$.

Public Key: $(\Delta_q, k, n, \mathfrak{p})$ or (Δ_q, k, n) , where

- $\Delta_q = q^2 \Delta_1$ with $\Delta_1 = p$;
- k and n are the bit lengths of $\lfloor \sqrt{\Delta_1}/4 \rfloor$ and $q - (\Delta_1/q)$, respectively;
- \mathfrak{p} is a randomly chosen \mathcal{O}_{Δ_q} -ideal with $[\mathfrak{p}] \in \ker(\hat{\Phi})$; inclusion of \mathfrak{p} in the public key is optional.

Here, p and q must be chosen so that R_{Δ_1} is small and R_{Δ_q} is large; details on how to select these primes will be provided in Section 5. If storage space for public keys is restricted, \mathfrak{p} need not be included in the public key. Instead, a different ideal \mathfrak{p} with $[\mathfrak{p}] \in \ker(\hat{\Phi})$ can be generated for each encryption, at the expense of increased encryption time. In the case where \mathfrak{p} is included in the public key, it can be generated exactly as in the original NICE system. In Section 7, we describe an alternative method for finding \mathfrak{p} that does not require knowledge of q and Δ_1 and can hence be used by the encrypter.

Encryption: Messages are bit strings of bit length $k - t - u$, where t and u are fixed parameters explained below. To encrypt a message m :

1. Convert m to a string m' that uniquely determines m and contains a predetermined bit pattern of length u .
2. Embed m' into a primitive \mathcal{O}_{Δ_q} -ideal \mathfrak{m} prime to q with $N_{\Delta_q}(\mathfrak{m}) \leq 2^k$ in such a way that $N_{\Delta_q}(\mathfrak{m})$ uniquely determines m' .
3. Generate random $r \in_R \{1, 2, \dots, 2^{n-1}\}$.
4. If the public key does not include the ideal \mathfrak{p} , generate a random \mathcal{O}_{Δ_q} -ideal \mathfrak{p} with $[\mathfrak{p}] \in \ker(\hat{\Phi})$.
5. The ciphertext is a reduced \mathcal{O}_{Δ_q} -ideal $\mathfrak{c} = \rho_{\Delta_q}(\mathfrak{m}\mathfrak{p}^r)$.

Decryption: To decrypt a ciphertext \mathcal{O}_{Δ_q} -ideal \mathfrak{c} :

1. Compute $\mathfrak{C} = \phi^{-1}(\mathfrak{c})$.
2. Find the reduced ideal $\mathfrak{M} \in [\mathfrak{C}]$ such that $N_{\Delta_1}(\mathfrak{M})$ contains the predetermined bit pattern of length u of encryption step 1.
3. Extract m' from $N_{\Delta_1}(\mathfrak{M})$ and m from m' .

Since the decrypter knows q , he can once again efficiently compute \mathfrak{C} ; for details, see Section 7. As in the original NICE scheme, we see that $[\mathfrak{M}] = [\phi^{-1}(\mathfrak{m})]$. Unfortunately, we can no longer conclude from this that $\mathfrak{M} = \phi^{-1}(m)$, only that both are two among many reduced ideals prime to q in the same \mathcal{O}_{Δ_1} -class. This is the reason why in contrast to encryption step 1 of NICE, the embedding of a message m into an \mathcal{O}_{Δ_q} -ideal \mathfrak{m} in REAL-NICE requires two steps. To ensure that $\mathfrak{M} = \phi^{-1}(\mathfrak{m})$ does in fact hold, m is endowed with a predetermined public bit pattern of length u to obtain m' . We argue below that this forces $\mathfrak{M} = \phi^{-1}(\mathfrak{m})$ with high probability, so $N_{\Delta_1}(\mathfrak{M}) = N_{\Delta_q}(\mathfrak{m})$ uniquely determines m' by encryption step 2, and hence m by encryption step 1.

More exactly, to perform encryption steps 1 and 2, one first selects the parameters t and u ; we explain below how large t and u should be chosen. The plaintext needs to be divided into blocks of bit length $k - t - u$. To embed such a block m into a reduced \mathcal{O}_{Δ_q} -ideal \mathfrak{m} prime to q , one does the following:

1. Set $m' = m + 2^{k-t}$, obtaining an integer m' of bit length $k - t$ whose u high order bits are $100 \cdots 000$.
2. Set $\overline{m'} = m'2^t$, obtaining an integer $\overline{m'}$ of bit length k whose u high order bits are $100 \cdots 000$ and whose t low order bits are all 0.
3. Find the smallest prime l exceeding $\overline{m'}$ such that $(\Delta_q/l) = 1$.
4. Set \mathfrak{m} to be the \mathcal{O}_{Δ_q} -ideal of norm l .

The ideal \mathfrak{m} is found exactly as in the NICE embedding procedure, and provided that $\mathfrak{M} = \phi^{-1}(\mathfrak{m})$, m' can again be extracted from $N_{\Delta_1}(\mathfrak{M}) = N_{\Delta_q}(\mathfrak{m}) = l$ with high probability by truncating the high order k bits from l . Then m is obtained from m' by simply discarding the u high order bits of m' .

Before we argue that, with high probability, the class of $[\mathfrak{C}]$ contains only one ideal whose norm contains our specified bit pattern (namely the ideal $\mathfrak{M} = \phi^{-1}(\mathfrak{m})$ of norm l), we explain how to find this ideal. In order to perform decryption step 2, the decrypter needs to traverse the set of reduced \mathcal{O}_{Δ_1} -ideals in the class of \mathfrak{C} to locate \mathfrak{M} . This is accomplished by applying repeated baby steps as described in Section 7, starting with the \mathcal{O}_{Δ_1} -ideal $\mathfrak{C} \in [\mathfrak{M}]$. Since Δ_1 was chosen so that the class of \mathfrak{C} contains very few reduced ideals, \mathfrak{M} can be found efficiently. After each baby step, the decrypter performs a simple X-OR on the u high order bits of the ideal norm and the string $100 \cdots 000$, checking whether or not the resulting string consists of all 0's.

The decryption procedure will work with high probability under two conditions. Firstly, just as in NICE, the parameter t needs to be chosen as described in Section 3 to ensure that m' can be uniquely determined from \mathfrak{M} . We already saw that this succeeds with probability at least $P_t = 1 - 2^{-2^t/k}$. Secondly, u must be chosen large enough so that with high probability, the \mathcal{O}_{Δ_1} -class of \mathfrak{C} contains

only one reduced \mathcal{O}_{Δ_1} -ideal \mathfrak{M} such that the u high order bits of $N_{\Delta_1}(\mathfrak{M})$ are $100 \cdots 000$. Using the analysis on p. 53, of [19], this probability is expected to be bounded below by $P_u = (1 - 2^{-u})^N$, where N is an upper bound on the number of reduced ideals in any class of $Cl(\mathcal{O}_{\Delta_1})$. We will see in Section 5 that Δ_1 can be chosen so that such an upper bound is of the form $\kappa_1 \log(\Delta_1)$ for some explicitly computable constant κ_1 .

5 Choice of Parameters

The parameters for REAL-NICE clearly need to be selected with care to ensure both efficiency and security. As explained in Section 4, p and q must be chosen to satisfy the following conditions:

- R_{Δ_q} must be large enough to ensure a sufficiently large number of reduced ideals in any \mathcal{O}_{Δ_q} -ideal class, thus rendering exhaustive search through any cycle of reduced \mathcal{O}_{Δ_q} -ideals infeasible.
- R_{Δ_1} must be small enough to ensure a sufficiently small number of reduced ideals in any \mathcal{O}_{Δ_1} -ideal class, thus rendering exhaustive search through any cycle of reduced \mathcal{O}_{Δ_1} -ideals efficient.

We proceed in two steps. First, we explain how to ensure that the ratio $R_{\Delta_q}/R_{\Delta_1}$ is of order of magnitude q with high probability. Then we present a means of guaranteeing that R_{Δ_1} is small, i.e. bounded by a polynomial in $\log(\Delta_1)$.

The *unit index* of \mathcal{O}_{Δ_q} is the group index $[\mathcal{O}_{\Delta_1}^* : \mathcal{O}_{\Delta_q}^*]$, i.e. the smallest positive integer i such that $\epsilon_{\Delta_1}^i = \epsilon_{\Delta_q}$, or equivalently, $R_{\Delta_q} = iR_{\Delta_1}$. By (2.3), i divides $q - (\Delta_1/q)$, so forcing i to be large is another reason why $q - (\Delta_1/q)$ should be almost prime. Specifically, Theorem 5.8, p. 58, of [19] states that if $q - (\Delta_1/q) = Ld$ where L is a large prime and $d \leq \log(\Delta_1)^\kappa$ for some positive constant κ , then the probability that $i < L$ is bounded above by $\log(\Delta_1)^{2\kappa}/(\sqrt{\Delta_1} - 1)$. In other words, $i = R_{\Delta_q}/R_{\Delta_1} \geq L$ with overwhelming probability.

To verify that $i \geq L$ does in fact hold, it suffices to check that i does not divide d , i.e. that $\epsilon_{\Delta_1}^d \neq \epsilon_{\Delta_q}$. Suppose that R_{Δ_1} is sufficiently small so that $\epsilon_{\Delta_1} = U_1 + V_1\sqrt{\Delta_1}$ is computable; ϵ_{Δ_1} can be obtained from R_{Δ_1} using for example Algorithm 4.2 of [2]. Then any power $\epsilon_{\Delta_1}^j = U_j + V_j\sqrt{\Delta_1}$ can be efficiently evaluated using Lucas function arithmetic on U_j and V_j analogous to binary exponentiation; see Chapter 4, pp. 69-95, of [20].

Next, we illustrate how to choose Δ_1 so that R_{Δ_1} is small. In general, the regulator R_{Δ_f} of any real quadratic order \mathcal{O}_{Δ_f} is of magnitude $\sqrt{\Delta_f}$ which is far too large for our purposes. One possibility is to choose D to be a *Schinzel sleeper* [15], i.e. a positive squarefree integer of the form $D = D(x) = a^2x^2 + 2bx + c$ with $a, b, c, x \in \mathbb{Z}$, $a \neq 0$, and $b^2 - a^2c$ dividing $4\gcd(a^2, b)^2$. Schinzel sleepers were analyzed in detail in [3]; here, the regulator R_{Δ_1} is of order $\log(\Delta_1)$. More exactly, if a, b, c, x are chosen so that $\gcd(a^2, 2b, c)$ is squarefree and $D \equiv 0$ or $1 \pmod{4}$ (so $\Delta_1 = D$), then by Theorem 5.4, p. 52, of [19], the number of reduced \mathcal{O}_{Δ_1} -ideals in any class of $Cl(\mathcal{O}_{\Delta_1})$ is bounded above by $\kappa_1 \log(\Delta_1)$ for an explicitly computable constant κ_1 that depends only on a and b .

Finally, the fixed bit pattern in any message is again critical in defending REAL-NICE against the same chosen ciphertext attack [7] that was already mentioned in Section 3. This attack can be detected with the same probability with which the cipher ideal can be successfully decrypted. Therefore, it is suggested to choose $t + u$ according to first row of Table 1, keeping the probability of successful decryption via the chosen ciphertext attack consistent with the probability of success of any other known attack on REAL-NICE.

6 Security

Although the security of REAL-NICE is based on a different mathematical problem than NICE, namely locating an \mathcal{O}_{Δ_q} -ideal within the cycle of reduced ideals in its own ideal class, as opposed to locating it in its own coset relative to $\ker(\hat{\phi})$, the same security considerations apply. Assuming a passive adversary, both systems can be broken if and only if an adversary can efficiently compute images of \mathcal{O}_{Δ_q} -ideals under the map ϕ^{-1} of (2.1) without knowledge of the trapdoor information q , a task that is provably equivalent to factoring in random polynomial time. More exactly, according to Theorem 2.1, pp. 13-14, of [18]:

Theorem 2. *Let $\Delta_1 \in \mathbb{N}$ be a fundamental discriminant and $\Delta_q = q^2\Delta_1$ with q prime. Assume that there exists an algorithm \mathbf{A} that computes for any primitive ideal $\mathfrak{a} \in \mathcal{I}(\mathcal{O}_{\Delta_q}, q)$ the primitive ideal $\mathfrak{A} = \phi^{-1}(\mathfrak{a}) \in \mathcal{I}(\mathcal{O}_{\Delta_1}, q)$ without knowledge of the conductor q of \mathcal{O}_{Δ_q} . By using the algorithm \mathbf{A} as an oracle, Δ_q can be factored in random polynomial time. The number of required queries to the oracle is polynomially bounded in $\log(\Delta_q)$.*

Hence, as with NICE, p and q must be chosen sufficiently large to render the factorization of Δ_q infeasible. Again, it is highly unlikely that knowledge of the public information would compromise the security of REAL-NICE. In addition, the specified bit pattern in the norm of the message ideal \mathfrak{m} protects against the chosen ciphertext attack of [7]; once again, the length of this bit pattern should be chosen equal to the symmetric key size as specified in Table 1 to render this attack as expensive as any other known attack. As REAL-NICE is so similar to NICE, it should also be possible to adapt the methods of [1] to obtain IND-CCA2 security.

The fact that $\Delta_1 = a^2x^2 + 2bx + c$ is chosen to be a Schinzel sleeper requires further analysis. It is recommended that the values a, b, c, x are kept secret and discarded after computing Δ_1 . Care must also be taken how to select x in the Schinzel sleeper. Put $A = qa$, $B = q^2b$, and suppose $B = SA + R$ with $0 \leq R < A$ (note that A, B, S, R are all unknown). Then by Theorem 4.1 of [3], the fraction A/R appears among the first $\kappa_2 \log(A)$ convergents of the continued fraction expansion of $\sqrt{\Delta_q}$ for some explicit positive constant κ_2 , so there are only polynomially many possibilities for this fraction. If we find A/R and write it in lowest terms, i.e. $A/R = U/V$ with $\gcd(U, V) = 1$, then $q = \gcd(\Delta_q, U \lfloor \sqrt{\Delta_q} \rfloor + V)$ if x is sufficiently large, so Δ_q is factored. This factoring attack can be avoided if x is chosen sufficiently small, but at the same time large enough to guarantee

sufficiently large Δ_1 . More exactly, by Corollary 6.5, p. 72, of [19], it is sufficient to choose $x < 2^{-3w-1}q - 2^w$ where $a, b \leq 2^w$.

Suppose we wish to generate parameters Δ_1 and q of bit length s . If we choose a and b of some bit length $w \leq s/4 - 1$ and x of bit length $s/2 - w$, then ax has bit length $s/2$, $2bx$ has bit length $s/2 + 1$, and the condition $b^2 - a^2c$ divides $4 \gcd(a^2, b)^2$ implies $|c| \leq 5b^2 < 2^{2w+3} \leq 2^{s/2+1}$, so $|c|$ has bit length at most $s/2 + 1$. Thus, $(ax)^2$ is the dominant term in $D(x)$, which then has bit length s . Since $q > 2^{s-1}$, it suffices to choose $x < 2^{s-3w-2} - 2^w$, so to obtain x of bit length at least $s/2 - w$, we require that $2^{s/2-w} < 2^{s-3w-2} - 2^w$. This is easily verified to always hold if $w \leq s/4 - 1$.

We also need to ensure that there are sufficiently many primes of desired size that occur as values of Schinzel sleepers. Let $\pi_F(n)$ denote the number of primes assumed by the polynomial $F(x) = ax^2 + bx + c$ for $0 \leq x \leq n$, with $a, b, c \in \mathbb{Z}$, $a > 0$, and $a + b, c$ not both even. The well-known Hardy-Littlewood conjecture [5] states in essence that $\pi_F(n) \sim \kappa_F n / \log(n)$, where κ_F is an explicitly computable constant than depends only on F . Under the assumption that prime values assumed by Schinzel polynomials behave similarly to those assumed by arbitrary quadratic polynomials, we conclude that the number of primes produced by Schinzel polynomials is large enough to render an exhaustive search for Δ_q infeasible. However, further study of this question is warranted.

Finally, we need to make sure that there are sufficiently many reduced \mathcal{O}_{Δ_q} -ideals of the form $\mathfrak{c} = \rho_{\Delta_q}(\mathfrak{m}\mathfrak{p}^r)$ to ensure that cipher ideals cannot be found via exhaustive search. We already saw how to guarantee a large ratio $R_{\Delta_q}/R_{\Delta_1} = Ld$ where L is a large prime and $d \leq \log(\Delta_1)^\kappa$ for some positive constant κ . This ensures a large number of reduced ideals in each \mathcal{O}_{Δ_q} -ideal class. For any $B \in \mathbb{N}$, any \mathcal{O}_{Δ_q} -ideal \mathfrak{p} with $[\mathfrak{p}] \in \ker(\hat{\Phi})$, and any reduced \mathcal{O}_{Δ_q} -ideal \mathfrak{m} , consider the set of possible cipher ideals $\mathcal{C}_B = \{\rho_{\Delta_q}(\mathfrak{m}\mathfrak{p}^r) \mid 1 \leq r \leq B\}$. Then a sufficiently large choice of a generator $\alpha \in \mathcal{O}_{\Delta_q}$ of \mathfrak{p} ensures that all the ideals in \mathcal{C}_B are distinct. More exactly, according to Theorem 6.8, p. 81, of [19], if we choose α so that $\log(\alpha) \in I$ where

$$I =](b + 1) \log(4\Delta_q) + \log(2), \frac{L \log(\Delta_1)}{2^{b+2}} - b \log(2)] \tag{6.4}$$

and b is the bit length of B , then the set \mathcal{C}_B has cardinality B .

Table 2 contains upper and lower bounds on $\log(\alpha)$ depending on the required level of security and the constant κ . The notation (v, w) in the column headers means that the set \mathcal{C}_B contains at least 2^v different \mathcal{O}_{Δ_q} -ideals, where Δ_q has bit length w . The columns “min” and “max” denote lower and upper bounds on the bit length of $\log(\alpha)$. The data show that it is feasible to choose α such that \mathcal{C}_B is sufficiently large to satisfy the NIST security requirements.

7 Ideal Arithmetic and Algorithms

We review basic ideal arithmetic involving \mathbb{Z} -bases and provide the algorithms that are required in the REAL-NICE cryptosystem. See also [21,9] for details.

Table 2. Bounds the size of $\log \alpha$ depending on the level of security

	(80, 1024)		(112, 2048)		(128, 3092)		(192, 8192)		(256, 15360)	
κ	min	max	min	max	min	max	min	max	min	max
1	17	259	18	568	20	1235	21	2536	22	4861
5	17	223	18	528	20	1191	21	2488	22	4809
10	17	178	18	478	20	1136	21	2428	22	4744
15	17	133	18	428	20	1081	21	2368	22	4679
20	17	88	18	378	20	1026	21	2308	22	4614

Let $\mathcal{K} = \mathbb{Q}(\sqrt{D})$ be a (real or imaginary) quadratic field and \mathcal{O}_{Δ_f} an order in \mathcal{K} of conductor f . Set $\sigma = 1$ if $D \equiv 2, 3 \pmod{4}$ and $\sigma = 2$ if $D \equiv 1 \pmod{4}$, so $\Delta_1 = (2/\sigma)^2 D \equiv \sigma - 1 \pmod{4}$. Then every integral \mathcal{O}_{Δ_f} -ideal \mathfrak{a} is a \mathbb{Z} -module of the form

$$\mathfrak{a} = S \left(\frac{Q}{\sigma} \mathbb{Z} + \frac{P + f\sqrt{D}}{\sigma} \mathbb{Z} \right),$$

where $S, Q \in \mathbb{N}$, $P \in \mathbb{Z}$, σ divides Q , σQ divides $f^2 D - P^2$, and $\gcd(Q, 2P, (f^2 D - P^2)/Q) = \sigma$. Here, Q and S are unique and P is unique modulo Q , so we write $\mathfrak{a} = S(Q, P)$ for brevity. We have $N_{\Delta_f}(\mathfrak{a}) = S^2 Q/\sigma$. The ideal \mathfrak{a} is primitive if and only if $S = 1$, in which case we simply write $\mathfrak{a} = (Q, P)$.

Suppose now that $D > 0$, so \mathcal{K} is a real quadratic field. Recall that any \mathcal{O}_{Δ_f} -ideal class contains a finite number of reduced ideals. A *baby step* moves from one such ideal to the next. More exactly, if $\mathfrak{a}_i = (Q_i, P_i)$ is a reduced \mathcal{O}_{Δ_f} -ideal, then a reduced \mathcal{O}_{Δ_f} -ideal $\mathfrak{a}_{i+1} = (Q_{i+1}, P_{i+1})$ in the \mathcal{O}_{Δ_f} -ideal class of \mathfrak{a}_i can be obtained using the formulas

$$q_i = \left\lfloor \frac{P_i + \sqrt{D}}{Q_i} \right\rfloor, \quad P_{i+1} = q_i Q_i - P_i, \quad Q_{i+1} = \frac{f^2 D - P_{i+1}^2}{Q_i}. \tag{7.5}$$

Note that $f^2 D = (\sigma/2)^2 \Delta_f$, so f need not be known here. Baby steps applied to any reduced \mathcal{O}_{Δ_f} -ideal \mathfrak{a} produce the entire cycle of reduced ideals in the \mathcal{O}_{Δ_f} -ideal class of \mathfrak{a} . In practice, one uses a more efficient version of (7.5) that avoids the division in the expression for Q_{i+1} ; see for example Algorithm 1 of [21].

We now give details on how to perform the different encryption and decryption steps, beginning with a method for finding a \mathbb{Z} -basis (Q, P) of the message ideal \mathfrak{m} of prime norm l as required in encryption step 2 of REAL-NICE. Set $Q = 2l$, and let P' be a square root of Δ_q modulo l with $0 < P' < l$. Such a square root exists since $(\Delta_q/l) = 1$ and can be found using standard probabilistic methods in at most an expected $O(\log(l)^3)$ bit operations. Now put $P = P'$ if P is odd and $P = l - P'$ if P is even. Then it is not hard to verify that $\mathfrak{m} = (Q, P)$ is a primitive \mathcal{O}_{Δ_q} -ideal of norm $Q/2 = l$.

Given \mathbb{Z} -bases of two reduced \mathcal{O}_{Δ_f} -ideals $\mathfrak{a}, \mathfrak{b}$, it is well-known how to compute a \mathbb{Z} -basis of a reduced \mathcal{O}_{Δ_f} -ideal $\rho_{\Delta_f}(\mathfrak{a}\mathfrak{b})$ in the class of the (generally non-reduced and possibly not even primitive) product ideal $\mathfrak{a}\mathfrak{b}$ in $O(\log(D)^2)$

bit operations. This operation is called a *giant step*. Five different ways for effecting a giant step were described and compared in [9]; the most efficient one is Algorithm 6.8 on p. 111 (the NUCOMP algorithm). This method can be employed to compute the cipher ideal \mathfrak{c} in REAL-NICE encryption step 5.

An ideal $\mathfrak{p} \in \ker(\hat{\phi})$ as required in encryption step 4 can be determined during encryption or as part of the public key as follows. Generate a random element $x \in I$, with I as given in (6.4), and use Algorithm 5.2 of [14] to find the \mathbb{Z} -basis of a reduced principal \mathcal{O}_{Δ_q} -ideal \mathfrak{p} that has a generator $\alpha \in \mathcal{O}_{\Delta_q}$ with $\log_2(\alpha) \approx x/\log(2)$, so $\log(\alpha) \approx x$. This algorithm is essentially repeated squaring using giant steps and requires $O(\log(x)\log(\Delta_q)^2)$ bit operations. Since \mathfrak{p} is principal and ϕ^{-1} preserves principality, $\phi^{-1}(\mathfrak{p})$ is a principal \mathcal{O}_{Δ_1} -ideal, so $[\mathfrak{p}] \in \ker(\hat{\phi})$.

In decryption step 1, the user needs to find the image \mathfrak{C} of the cipher ideal \mathfrak{c} under ϕ^{-1} . The functions ϕ^{-1} and ϕ can be efficiently computed if the conductor f is known. We briefly recall the procedures here; for details, see [6,13] as well as pp. 14 and 28 of [19]. Let $\mathfrak{a} = (Q, P_{\Delta_f})$ be any primitive \mathcal{O}_{Δ_f} -ideal prime to f . Then $\mathfrak{A} = \phi^{-1}(\mathfrak{a}) = (Q, P_{\Delta_1})$ is a primitive \mathcal{O}_{Δ_1} -ideal prime to f , where $P_{\Delta_1} \equiv xP_{\Delta_f} + ybQ/2 \pmod{Q}$. Here, $x, y \in \mathbb{Z}$ are given by $xf + yQ/\sigma = 1$, and b is the parity of Δ_f ; note that if Q is odd, then $\sigma = 1$ and hence $b = 0$. Conversely, if $\mathfrak{A} = (Q, P_{\Delta_1})$ is a primitive \mathcal{O}_{Δ_1} -ideal prime to f , then $\mathfrak{a} = \phi(\mathfrak{A}) = (Q, P_{\Delta_f})$ with $P_{\Delta_f} \equiv fP_{\Delta_1} \pmod{Q}$ is a primitive \mathcal{O}_{Δ_f} -ideal prime to f .

8 Implementation and Run Times

We implemented prototypes of both NICE and REAL-NICE in C++ using GMP and the NTL library for large integer arithmetic [16]. Our numerical data were generated on an Athlon XP 2000+ with 512 MB RAM under the Linux Mandrake 9.1 operating system. In addition, we felt that a comparison to the RSA cryptosystem would be of interest, since the security of RSA also depends on integer factorization and, because it is so widely used in practice, highly-optimized implementations are readily available. We therefore determined run times for RSA using the open source implementation of OpenSSL. Note that our current implementations of NICE and REAL-NICE are first prototypes, whereas the RSA implementation in OpenSSL is highly optimized. Thus, our numerical results are somewhat skewed in favour of RSA.

We used the same parameter sizes for our RSA moduli and non-fundamental discriminants $\Delta_q = q^2\Delta_1$, with q and Δ_1 of approximately equal size, as they give the same level of security. We chose parameter sizes corresponding to the NIST recommended levels of security equivalent to block ciphers with keys of 80, 112, 128, 192, and 256 bits, as specified in Table 1. As public key cryptosystems are usually used for secure key exchange, the message lengths used in our set-up corresponded to these key sizes.

Both NICE and REAL-NICE require selecting a suitable fundamental discriminant Δ_1 . For NICE, we simply chose $\Delta_1 = -p$ where p is a prime with $p \equiv 3 \pmod{4}$. In REAL-NICE, we chose $\Delta_1 = p$ where $p \equiv 1 \pmod{4}$ and

p is a Schinzel sleeper as described in Section 5. To find such a prime p of bit length s , we began by choosing random positive integers a and b of bit length $w = 12$; this length easily satisfies the requirement $w \leq s/4 - 1$ of Section 6 for all five NIST [12] security levels with s as given in the second row of Table 1. Then we attempted to determine an integer c such that $b^2 - a^2c$ divides g^2 with $g = 2 \gcd(a^2, b)$. To find c , we searched the interval $S = [(b^2 - g^2)/a^2, (b^2 + g^2)/a^2]$ for an integer c satisfying the above divisibility condition. Note that if a and b are randomly generated, then $g = 1$ with high probability, leaving only a very small search interval S . Moreover, smaller values of a lead to a larger interval S . Hence, if for a given pair (a, b) , no suitable c value was found, we decreased a by 1 — rather than generating a new random value a — and conducted a new search for c . We repeated this procedure until either $a = 1$ or a suitable value of c was found; in the former case, we discarded a and b and started over.

Once a suitable triple (a, b, c) was obtained, we generated successive random integers x of bit length at least $s/2 - 12$ until a value x was found such that $ax^2 + 2bx + c$ is a prime congruent to 1 (mod 4). After a certain number of unsuccessful trials at a value of x , we discarded the triple (a, b, c) and started over with a new choice of a and b . This method worked very well in practice.

To find a conductor q such that $q - (\Delta_1/q)$ is guaranteed to have a large prime factor, we first generated a random prime L close to $|\Delta_1|$ and checked exhaustively whether $q = jL + 1$ is prime and $(\Delta_1/q) = -1$ for $j = 2, 4, 6, \dots$. If no prime was found for j up to some predetermined bound M , we discarded L and repeated the same procedure until a prime l with the desired properties was obtained.

For encryption under NICE and REAL-NICE, messages were embedded into an \mathcal{O}_{Δ_q} -ideal of prime norm l such that the binary representation of l contained a fixed bit pattern of length $b_{\Delta_q} \in \{80, 112, 128, 192, 256\}$ corresponding to the level of security that was chosen for Δ_q ; $b_{\Delta_q} = t$ in NICE and $b_{\Delta_q} = t + u$ in REAL-NICE. In addition, the 20 low order bits of l were set so that $(\Delta_q/l) = 1$. Consequently, messages were bit strings of length $k - b_{\Delta_q} - 20$.

Table 3 gives the average run times for NICE, REAL-NICE, and RSA for various parameter sizes. The run times were obtained by encrypting and decrypting 1000 randomly generated messages for each discriminant size. In addition to the timings for encryption, decryption and message embedding, Table 3 lists the minimal, the maximal and the average number of baby steps that were required to locate the \mathcal{O}_{Δ_1} -ideal \mathfrak{M} during decryption with REAL-NICE.

Our numerical results show that NICE out-performs REAL-NICE for both encryption and decryption. This is not surprising. Recall that in REAL-NICE, the ideal \mathfrak{p} with $[\mathfrak{p}] \in \ker(\hat{\phi})$ is not included in the public key, resulting in shorter keys. This is done at the expense of a considerable increase in encryption time due to the need for generating a new random ideal \mathfrak{p} for each encryption. We also expect decryption times of REAL-NICE to be slower than those of NICE, due to the extra search through the cycle of reduced ideals in the class of the \mathcal{O}_{Δ_1} -ideal $\mathfrak{C} = \phi^{-1}(\mathfrak{c})$, where \mathfrak{c} is the cipher ideal. In fact, decryption showed the most significant difference in performance between NICE and REAL-NICE. When

Table 3. Average Run Times for NICE, REAL-NICE, and RSA

size(Δ_q)	1024	2048	3072	8192	15360
message length	80	112	128	192	256
block length	180	244	276	404	532
NICE					
encryption	0.02139s	0.06994s	0.12659s	0.68824s	2.35274s
decryption	0.00033s	0.00082s	0.00099s	0.00312s	0.00729s
embedding	0.00467s	0.01152s	0.01550s	0.04257s	0.08860s
REAL-NICE					
encryption	0.03532s	0.09944s	0.18096s	0.96830s	3.28507s
decryption	0.00210s	0.00468s	0.00757s	0.02811s	0.07735s
embedding	0.00531s	0.01152s	0.01547s	0.04289s	0.09770s
min. number of baby steps	1	1	1	2	2
max. number of baby steps	127	181	193	271	355
avg. number of baby steps	58.345	92.801	121.107	204.056	281.438
RSA					
encryption	0.0074s	0.0081s	0.0090s	0.0173s	0.0499s
decryption	0.0127s	0.0334s	0.0931s	1.1188s	7.8377s

considering the overall performance, NICE is up to 1.61 faster than REAL-NICE. However, we note that encryption in REAL-NICE can be replaced by a technique called *infrastructure multiple side-step* (IMS) encryption that could potentially make REAL-NICE competitive to NICE; a similar idea was used with considerable success in cryptographic protocols using real hyperelliptic curves [8], and is explained in the next section.

As expected, both NICE and REAL-NICE decryption significantly outperform RSA for all security levels. It is also noteworthy that when considering the overall performance, both NICE and REAL-NICE are faster than RSA for the two highest levels of security. This is surprising as our NICE and REAL-NICE implementations are first prototypes, whereas the implementation of RSA in the OpenSSL package is considered to be highly optimized.

9 Conclusion and Further Work

There exists a modified version of RSA due to Takagi [17] that would perhaps be more appropriate for comparison with NICE and REAL-NICE. Takagi's cryptosystem relies on the difficulty of factoring integers of the form $p^k q$ (similar to NICE and REAL-NICE) and has faster decryption than RSA. When using $k = 2$, Takagi reports decryption times that are three times faster than decryption using Chinese remaindering with a 768-bit modulus. Our main goal was to compare NICE and REAL-NICE with a highly-optimized implementation of the most widely-used factoring-based cryptosystem (namely RSA), but a comparison with Takagi's cryptosystem would clearly be of interest as well.

While the performance difference between NICE and REAL-NICE is at first glance disappointing, a method referred to as *infrastructure multiple side-step* (IMS) encryption can speed up REAL-NICE encryption time considerably, making the system potentially competitive with NICE. IMS is explained in detail in Section 7.2.1 of [19]. In both NICE and REAL-NICE, the \mathcal{O}_{Δ_q} -ideal $\rho_{\Delta_q}(\mathfrak{p}^r)$ used to obtain the cipher ideal \mathfrak{c} is evaluated using a standard binary exponentiation technique involving giant steps. That is, a square corresponds to a giant step of the form $\rho_{\Delta_q}(\mathfrak{a}^2)$, and a multiply to the giant step $\rho_{\Delta_q}(\mathfrak{a}\mathfrak{p})$, where \mathfrak{a} is the intermediate ideal ($\mathfrak{a} = \rho_{\Delta_q}(\mathfrak{p}^r)$ at the end). In IMS encryption, no random exponent needs to be generated. Instead, a fixed number of square giant steps is chosen, and each square giant step is followed by a certain random number of baby steps (*multiple side steps*) in the cycle of reduced ideals (also referred to as the *infrastructure*) in the \mathcal{O}_{Δ_q} -class of the message ideal \mathfrak{m} .

The complexity of a baby step is linear in $O(\log(\Delta_q))$ in terms of bit operations, whereas a giant step has quadratic complexity. Thus, if the number of square giant steps corresponds to the bit length of r , and the number of baby steps after each square & reduce operation is not too large, this results in a significant speed-up in encryption time. On the other hand, if the number of squarings or the number of side steps is too small, this may significantly decrease the number of possible values that the cipher ideal \mathfrak{c} can take on, thereby rendering exhaustive search for \mathfrak{c} potentially feasible. Preliminary numerical data in Section 7.3 of [19] showed that an IMS-prototype of REAL-NICE using even the most conservative security analysis outperformed the original REAL-NICE scheme. It is conceivable that the IMS parameters could be chosen to lead to significantly faster encryption times, while still ensuring the same level of security. Under these circumstances, IMS-REAL-NICE could be competitive to, or even outperform, NICE. This would make IMS-REAL-NICE potentially attractive in situations where fast decryption time is essential (e.g. for fast signature generation) and space is too restricted to hold the larger NICE keys. Clearly, the subject of IMS encryption requires further exploration.

The questions of whether there are sufficiently many prime Schinzel sleepers of a given bit length, and whether choosing Δ_1 to be a Schinzel sleeper presents a security risk, warrant further study. We also point out that it should be possible to adapt the IND-CCA2 secure version of NICE to REAL-NICE in order to provide the same level of security. These and other questions are the subject of future research.

References

1. Buchmann, J., Sakurai, K., Takagi, T.: An IND-CCA2 public-key cryptosystem with fast decryption. In: Kim, K.-c. (ed.) ICISC 2001. LNCS, vol. 2288, pp. 51–71. Springer, Heidelberg (2002)
2. Buchmann, J., Thiel, C., Williams, H.: Short representation of quadratic integers. In: Computational Algebra and Number Theory (Sydney, 1992). Math. Appl., vol. 325, pp. 159–185. Kluwer, Dordrecht (1995)

3. Cheng, K.H.F., Williams, H.C.: Some results concerning certain periodic continued fractions. *Acta Arith.* 117, 247–264 (2005)
4. Cox, D.A.: Primes of the Form $x^2 + ny^2$. John Wiley & Sons, Inc, New York (1989)
5. Hardy, G.H., Littlewood, J.E.: Partitio numerorum III: On the expression of a number as a sum of primes. *Acta Math.* 44, 1–70 (1923)
6. Hühnlein, D., Jacobson Jr., M.J., Paulus, S., Takagi, T.: A cryptosystem based on non-maximal imaginary quadratic orders with fast decryption. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 294–307. Springer, Heidelberg (1998)
7. Jaulmes, E., Joux, A.: A NICE Cryptanalysis. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 382–391. Springer, Heidelberg (2000)
8. Jacobson Jr., M.J., Scheidler, R., Stein, A.: Cryptographic protocols on real hyperelliptic curves. *Adv. Math. Commun.* 1, 197–221 (2007)
9. Jacobson Jr., M.J., Sawilla, R.E., Williams, H.C.: Efficient Ideal Reduction in Quadratic Fields. *Internat. J. Math. Comput. Sci.* 1, 83–116 (2006)
10. Lenstra, A.K.: Unbelievable Security. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 67–86. Springer, Heidelberg (2001)
11. Mollin, R.A., Williams, H.C.: Computation of the class number of a real quadratic field. *Util. Math.* 41, 259–308 (1992)
12. National Institute of Standards and Technology (NIST), Recommendation for key management - part 1: General (revised). NIST Special Publication 800-57 (March 2007), http://csrc.nist.gov/groups/ST/toolkit/documents/SP800-57Part1_3-8-07.pdf
13. Paulus, S., Takagi, T.: A new public key cryptosystem over quadratic orders with quadratic decryption time. *J. Cryptology* 13, 263–272 (2000)
14. van der Poorten, A.J., te Riele, H.J.J., Williams, H.C.: Computer verification of the Ankeny-Artin-Chowla conjecture for all primes less than 100 000 000 000. *Math. Comp.* 70, 1311–1328 (2001)
15. Schinzel, A.: On some problems of the arithmetical theory of continued fractions. *Acta Arith.* 6, 393–413 (1961)
16. Shoup, V.: NTL: A Library for Doing Number Theory. Software (2001), Available at <http://www.shoup.net/ntl>
17. Takagi, T.: Fast RSA-type cryptosystem modulo p^kq . In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 318–326. Springer, Heidelberg (1998)
18. Takagi, T.: A New Public-Key Cryptosystems with Fast Decryption. PhD Thesis, Technische Universität Darmstadt (Germany) (2001)
19. Weimer, D.: An Adaptation of the NICE Cryptosystem to Real Quadratic Orders. Master's Thesis, Technische Universität Darmstadt (Germany) (2004), <http://www.cdc.informatik.tu-darmstadt.de/reports/reports/DanielWeimer.diplom.pdf>
20. Williams, H.C.: Édouard Lucas and Primality Testing. John Wiley & Sons, New York (1998)
21. Williams, H.C., Wunderlich, M.C.: On the parallel generation of the residues for the continued fraction factoring algorithm. *Math. Comp.* 48, 405–423 (1987)