Decision Oracles Are Equivalent to Matching Oracles

[Published in H. Imai and Y. Zheng, Eds., *Public-Key Cryptography*, vol. 1560 of *Lecture Notes in Computer Science*, pp. 276–289, Springer-Verlag, 1999.]

Helena Handschuh[1], Yiannis Tsiounis[2], and Moti Yung[3]

[1] Gemplus Card International – ENST, Paris, France
handschuh@gemplus.com, handschu@enst.fr
[2] GTE Laboratories, Inc., Waltham, MA., USA
ytsiounis@gte.com
[3] CertCo Inc., New York, NY., USA
moti@cs.columbia.edu

Abstract. One of the key directions in complexity theory which has also filtered through to cryptographic research, is the effort to classify related but seemingly distinct notions. Separation or reduction arguments are the basic means for this classification.

Continuing this direction we identify a class of problems, called "matching problems," which are related to the class of "decision problems." In many cases, these classes are neither trivially equivalent nor distinct. Briefly, a "decision" problem consists of one instance and a supposedly related image of this instance; the problem is to decide whether the instance and the image indeed satisfy the given predicate. In a "matching" problem two such pairs of instances-images are given, and the problem is to "match" or "distinguish" which image corresponds to which instance. Clearly the decision problem is more difficult, since given a "decision" oracle one can simply test each of the two images to be matched against an instance and solve the matching problem. Here we show that the opposite direction also holds, presuming that randomization of the input is possible, and that the matching oracle is successful in all but a negligible part of its input set.

We first apply our techniques to show equivalence between the matching Diffie-Hellman and the decision Diffie-Hellman problems which were both applied recently quite extensively. This is a constructive step towards examining the strength of the Diffie-Hellman related problems. Then we show that in cryptosystems which can be uniformly randomized, non-semantic security implies that there is an oracle that decides whether a given plaintext corresponds to a given ciphertext. In the process we provide a new characteristic of encryption functions, which we call "universal malleability."

Keywords. Diffie-Hellman variants, randomized reductions, uniform reductions, public-key encryption, homomorphic encryption functions (ElGamal, Goldwasser-Micali, Okamoto-Uchiyama, Naccache-Stern), random self-reducibility, decision problems, matching problems, universal malleability.

1 Introduction

Recently we have seen a number of constructions that are based on the difficulty of the decision Diffie-Hellman problem, ranging from ElGamal-based encryption [ElG85,Dam91,TY98,CS98], to electronic cash [FTY98], and to hash functions [Can97]. A few variations of this problem, called "matching Diffie-Hellman" have also appeared [FTY96,Oka95,CFT98], mainly in electronic cash systems. Our first goal in this paper is to investigate the relationships between these variations and the original problem. To this extent we conclude that the problems are equivalent; this can be seen either as a strengthening of the matching Diffie-Hellman assumptions, or as a weakening of the decision Diffie-Hellman assumption.

Since our reduction techniques for deriving this equivalence are general, they can be applied to other settings in order to transform matching oracles to decision oracles. One such setting is the setting of secure encryption, i.e., the concept of indistinguishability of encryptions. In this context we show, under a specific assumption about the encryption scheme, that distinguishability of encryptions allows us to decide whether a given plaintext corresponds to a given ciphertext. Loosely speaking, this direction enhances the relationship between indistinguishability and semantic security in the sense that it provides, even if only for a limited set of cryptosystems, a specific kind of information that can be retrieved about a ciphertext, if the encryption is not secure in the sense of indistinguishability.

In the course of defining the properties that we require from a cryptosystem that allows this "attack," we propose a new definition, that of *universally malleable* cryptosystems. Intuitively, these are encryption schemes in which, without knowledge of the secret key, one can randomize, independently, both the message and the ciphertext. Typically this property is derived from the random self-reducibility of some underlying problem. Examples of such systems are the ElGamal cryptosystem [ElG85], the Okamoto-Uchiyama factoring-based cryptosystem [OU98], the Naccache-Stern higher-order residue cryptosystem [NS98], and the Goldwasser-Micali quadratic-residue cryptosystem [GM84].

Finally, one can use our methodology to show equivalences between general decision and matching problems. However, the equivalence can be shown only when the "matching oracle" can operate on all but a negligible part of the problem set and when inputs to the oracle can be randomized; this is why the universal malleability is required for the case of encryption systems.

Organization: In section 2 we define the matching and decision Diffie-Hellman problems. We proceed to collapse the decision problem to the matching, i.e., prove equivalence, in section 3. In section 4 we apply our result to cryptosystems, and in section 5 we discuss additional variants of the matching Diffie-Hellman problem.

2 Matching Diffie-Hellman and Decision Diffie-Hellman

In this section we formally define the Matching Diffie-Hellman and the Decision Diffie-Hellman problems. We begin by defining the common setting.

Common setting. For security parameter n, primes P and Q are chosen such that $|P-1| = \delta + n$ for a specified constant δ, and $P = \gamma Q + 1$, for a specified integer γ. Then a unique subgroup G_Q of prime order Q of the multiplicative group Z_P^* and a generator g of G_Q are defined. All the subsequent calculations are performed mod P, except operations involving exponents, which are performed mod Q.

Definition 1. (Decision Diffie-Hellman Problem) *For security parameter n, P a prime with $|P-1| = \delta+n$ for a specified constant δ, for $g \in Z_P^*$ a generator of prime order $Q = (P-1)/\gamma$ for a specified integer γ and for uniformly chosen $a, b \in_R G_Q$, given $[g^a, g^b, y]$ output 0 if $y = g^{ab} \pmod P$ and 1 otherwise, with probability better than $\frac{1}{2} + \frac{1}{n^c}$ for some constant c for large enough n.*

The *decision Diffie-Hellman assumption* (DDH) states that it is infeasible for a p.p.t. adversary to solve the Decision Diffie-Hellman problem.

Definition 2. (Matching Diffie-Hellman Problem) *For security parameter n, for uniformly chosen $a_i, b_i \in_R G_Q$ $(i \in \{0,1\})$, P a prime with $|P-1| = \delta + n$ for a specified constant δ, and for $g \in Z_P^*$ a generator of prime order $Q = (P-1)/\gamma$ for a specified small integer γ, given $[g^{a_0}, g^{b_0}]$, $[g^{a_1}, g^{b_1}]$ and $g^{a_r b_r}$,$g^{a_{\bar{r}} b_{\bar{r}}}$, $r, \bar{r} \in_R \{1,0\}, r \oplus \bar{r} = 1$, find r with probability better than $\frac{1}{2} + \frac{1}{n^c}$ for some constant c for large enough n.*

The *matching Diffie-Hellman assumption* (MDH) states that it is infeasible for a p.p.t. adversary to solve the Matching Diffie-Hellman problem.

Clearly, the DDH problem is at least as hard as the MDH since via two calls to a decision oracle we can solve the matching problem. The goal of the next section is to show the equivalence of these two problems. Intuitively, the problem of mapping the right Diffie-Hellman triplets *together* seems related to deciding whether a given triplet is a *correct* Diffie-Hellman triplet or not. But it is not clear whether, and how, one can use the seemingly weaker *matching* oracle to solve the *decision* problem. Here we prove the reduction by giving an exact construction to achieve it. We only show one direction (matching oracle to decision oracle) since the converse is straightforward.

These results can be extended to the case where an adversary has to select which of two ciphertexts maps to which of two plaintexts (indistinguishability of encryptions), versus where she has to decide whether a given ciphertext is the encryption of a given plaintext. In other words, we show that indistinguishability of encryptions (and therefore semantic security) is equivalent to deciding whether a given ciphertext corresponds to a given plaintext. This however only holds under a specific assumption on the encryption scheme. Under this assumption, this is an extension of the notion of “matching” (distinguishability) of two ciphertext/plaintext pairs, as traditionally defined in [Gol93].

3 Matching Diffie-Hellman Is At Least as Hard as Decision Diffie-Hellman

In this section we show how an attacker, given an oracle that solves the MDH problem with probability non negligibly better than $\frac{1}{2}$ (random guessing), can decide whether a given triplet is a correct Diffie-Hellman triplet or not with probability non negligibly better than random guessing. We are dealing with the uniform case.

Theorem 1. *Assume that there exists a probabilistic polynomial time Turing Machine which given an instance of the Matching Diffie-Hellman Problem gives the correct answer with probability better than* $\frac{1}{2} + \frac{1}{n^c}$ *for some constant* c *for large enough* n*. Then, there exists a p.p.t. TM which, given an instance of the Decision Diffie-Hellman Problem, gives the correct answer with probability better than* $\frac{1}{2} + \frac{1}{n'^{c'}}$ *for some constant* c' *for large enough* n'*.*

Proof. The proof is constructive. We show the steps that an adversary needs to take so that given a decision Diffie-Hellman problem she can solve it using the matching Diffie-Hellman oracle. This involves two phases.

1. **Testing Phase.**
   In this phase the oracle's behavior on incorrect inputs is tested. We will show that the oracle distinguishes either between two correct triplets and a correct and an incorrect one, *or* between a correct and an incorrect triplet and two incorrect ones. This fact will be used in the next phase to help us decide on whether the candidate Diffie-Hellman triplet is correct or not.
   First observe that if the oracle is given two random (i.e., non Diffie-Hellman) triplets, it cannot guess the attacker's random coin tosses for r, simply because no information (in the Shannon sense) about r is encoded in the input to the oracle. Formally, assume that the attacker uniformly and independently selects $r \in_R \{0,1\}, a_i, b_i \ (i \in \{0,1\}), v, w \in_R G_Q$, and then uses the oracle to estimate the quantity:[1]

$$\left| \Pr[A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^v, g^w) = r] - \frac{1}{2} \right| ,$$

   where $v, w \not\equiv a_i b_i \pmod{Q}$, for $i \in \{0,1\}$. It is clear that the probability of the oracle in finding r better than random guessing is negligible, since r is chosen randomly and independently of v, w and no information about r is included in the oracle's input. For clarity, we assume that the attacker has a success rate less than $\frac{1}{2n^c}$, i.e., that the oracle is run sufficiently many (polynomial) times so that the accuracy is $\frac{1}{2n^c}$. So we have that

$$\left| \Pr[A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^v, g^w) = r] - \frac{1}{2} \right| \leq \frac{1}{2n^c} \tag{1}$$

[1] Note that the notation $A[x] = r$ is a shortcut of saying that the adversary produces the correct "match". Thus we implicitly assume that an answer of 0 means that the first pair, in this case g^{a_0}, g^{b_0}, matches with the first number (g^v); and vice-versa, an answer of 1 means that the first pair matches with the second number.

On the other hand, from the assumption on the power of the oracle, we know that

$$\Pr[A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^{a_r b_r}, g^{a_{\bar{r}} b_{\bar{r}}}) = r] - \frac{1}{2} > \frac{1}{n^c} \ . \qquad (2)$$

In other words, the difference of behavior between two random triplets and two correct triplets is non-negligible. In particular, we have the following:

Lemma 1. *For every $a_i, b_i, c_i, d_i, i \in \{0,1\}$, for uniformly and independently chosen $r \in_R \{0,1\}$, $v, w \in_R G_Q$, and for large enough n, it holds that:*[2]

$$\begin{aligned}\Delta([a,b,1,1],[c,d,0,0]) = | \ &Pr\ [A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^{a_r b_r}, g^{a_{\bar{r}} b_{\bar{r}}}) = r] - \\ &Pr\ [A([g^{c_0}, g^{d_0}], [g^{c_1}, g^{d_1}], g^{v}, g^{w}) = r] \ | > \frac{1}{2n^c}\end{aligned}$$

Proof. First, from equation (1) we have

$$-\left|\Pr[A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^{v}, g^{w}) = r] - \frac{1}{2}\right| \geq -\frac{1}{2n^c}$$

Proceeding to prove the claim, we use the above together with equation (2) to get:

$$\begin{aligned}\Delta([a,b,1,1],\ &[c,d,0,0]) \\ &= |\Pr[A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^{a_r b_r}, g^{a_{\bar{r}} b_{\bar{r}}}) = r] \ - \\ &\qquad \Pr[A([g^{c_0}, g^{d_0}], [g^{c_1}, g^{d_1}], g^{v}, g^{w}) = r] \ | \\ &= |(\Pr[A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^{a_r b_r}, g^{a_{\bar{r}} b_{\bar{r}}}) = r] - \tfrac{1}{2}) \ - \\ &\qquad (\Pr[A([g^{c_0}, g^{d_0}], [g^{c_1}, g^{d_1}], g^{v}, g^{w}) = r] - \tfrac{1}{2}) \ | \\ &\geq |\Pr[A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^{a_r b_r}, g^{a_{\bar{r}} b_{\bar{r}}}) = r] - \tfrac{1}{2}| \ - \\ &\qquad |\Pr[A([g^{c_0}, g^{d_0}], [g^{c_1}, g^{d_1}], g^{v}, g^{w}) = r] - \tfrac{1}{2} \ | \\ &> \frac{1}{n^c} - |\Pr[A([g^{c_0}, g^{d_0}], [g^{c_1}, g^{d_1}], g^{v}, g^{w}) = r] - \tfrac{1}{2} \ | \\ &\geq \frac{1}{n^c} - \frac{1}{2n^c} \geq \frac{1}{2n^c}\end{aligned}$$

Now we show how the actual testing phase proceeds. First the attacker uniformly selects $r \in_R \{0,1\}$, $v \in_R G_Q$ and estimates the difference

$$\begin{aligned}\Delta([a,b,1,1],[e,f,1,0]) = | \ &\Pr\ [A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^{a_r b_r}, g^{a_{\bar{r}} b_{\bar{r}}}) = r] - \\ &\Pr\ [A([g^{e_0}, g^{f_0}], [g^{e_1}, g^{f_1}], g^{x}, g^{y}) = r] \ | \ ,\end{aligned}$$

[2] The notation here is as follows: $[a, b, i, j]$ signifies that a pair of triplets is given, such that when i (resp. j) is 0 the first (resp. the second) triplet is incorrect, and when it is 1 the triplet is a correct D-H triplet.

where $x, y \in_R \{e_r f_r, v\}$. The estimate is given with accuracy $\frac{1}{16n^c}$. Now if the estimate is greater or equal to $\frac{3}{16n^c}$ then the actual difference is at least $\frac{3}{16n^c} - \frac{1}{16n^c} = \frac{1}{8n^c}$. In this case we will say that the attacker can distinguish between two correct triplets and one correct/one incorrect triplet. If the estimate on the other hand is less than $\frac{3}{16n^c}$ then the actual difference is less than $\frac{3}{16n^c} + \frac{1}{16n^c} = \frac{1}{4n^c}$. In this case we say that the attacker cannot distinguish.
Now we will show that if the attacker cannot distinguish as above, then it must be able to distinguish between one correct/one incorrect triplet and two incorrect triplets. Starting from lemma 1, we have (definitions of variables are similar as above; we omit details):

$$\begin{aligned}
\frac{1}{2n^c} &< \Delta([a,b,1,1],[c,d,0,0]) \\
&= |\ \Pr[A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^{a_r b_r}, g^{a_{\bar r} b_{\bar r}}) = r] - \\
&\quad \Pr[A([g^{c_0}, g^{d_0}], [g^{c_1}, g^{d_1}], g^{v}, g^{w}) = r]\ | \\
&= |\ \Pr[A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^{a_r b_r}, g^{a_{\bar r} b_{\bar r}}) = r] - \\
&\quad \Pr[A([g^{e_0}, g^{f_0}], [g^{e_1}, g^{f_1}], g^{x}, g^{y}) = r] + \\
&\quad \Pr[A([g^{e_0}, g^{f_0}], [g^{e_1}, g^{f_1}], g^{x}, g^{y}) = r] - \\
&\quad \Pr[A([g^{c_0}, g^{d_0}], [g^{c_1}, g^{d_1}], g^{v}, g^{w}) = r]\ | \\
&\leq |\ \Pr[A([g^{a_0}, g^{b_0}], [g^{a_1}, g^{b_1}], g^{a_r b_r}, g^{a_{\bar r} b_{\bar r}}) = r] - \\
&\quad \Pr[A([g^{e_0}, g^{f_0}], [g^{e_1}, g^{f_1}], g^{x}, g^{y}) = r]\ | + \\
&\quad |\ \Pr[A([g^{e_0}, g^{f_0}], [g^{e_1}, g^{f_1}], g^{x}, g^{y}) = r] - \\
&\quad \Pr[A([g^{c_0}, g^{d_0}], [g^{c_1}, g^{d_1}], g^{v}, g^{w}) = r]\ | \\
&= \Delta([a,b,1,1],[e,f,1,0]) + \Delta([e,f,1,0],[c,d,0,0])
\end{aligned}$$

Thus, for uniformly chosen $e_i, f_i, i \in \{0,1\}$, i.e., $\Pr[(e_i, f_i)] = \frac{1}{|G_Q{}^2|}$, and for j enumerating all possible pairs, we have:

$$\begin{gathered}
\Sigma_j\ [\Delta([a,b,1,1],[e,f,1,0]) + \Delta([e,f,1,0],[c,d,0,0])] > \Sigma_j \frac{1}{2n^c} \iff \\
\Sigma_j\ \Delta([a,b,1,1],[e,f,1,0]) + \Sigma_j\ \Delta([e,f,1,0],[c,d,0,0])] > |G_Q{}^2| \frac{1}{2n^c} \iff \\
\Sigma_j\ \frac{\Pr[(e_i, f_i)]}{\Pr[(e_i, f_i)]} \Delta([a,b,1,1],[e,f,1,0]) + \\
\Sigma_j\ \frac{\Pr[(e_i, f_i)]}{\Pr[(e_i, f_i)]} \Delta([e,f,1,0],[c,d,0,0])] > |G_Q{}^2| \frac{1}{2n^c} \iff \\
|G_Q{}^2| \Sigma_j\ \Pr[(e_i, f_i)] \Delta([a,b,1,1],[e,f,1,0]) + \\
|G_Q{}^2| \Sigma_j\ \Pr[(e_i, f_i)] \Delta([e,f,1,0],[c,d,0,0])] > |G_Q{}^2| \frac{1}{2n^c} \iff \\
\Sigma_j\ \Pr[(e_i, f_i)] \Delta([a,b,1,1],[e,f,1,0]) + \\
\Sigma_j\ \Pr[(e_i, f_i)] \Delta([e,f,1,0],[c,d,0,0])] > \frac{1}{2n^c} \iff \\
E[\Delta([a,b,1,1],[e,f,1,0])] + E[\Delta([e,f,1,0],[c,d,0,0])] > \frac{1}{2n^c}\ ,
\end{gathered}$$

where the expectancy is taken over the choice of the triplets (e_i, f_i).
Therefore, if $E[\Delta([a,b,1,1],[e,f,1,0])] < \frac{1}{4n^c}$ then we have $E[\Delta([e,f,1,0],[c,d,0,0])] > \frac{1}{4n^c}$.
In summary, the oracle can be used to distinguish either between two correct triplets and one correct/one incorrect triplet, *or* between one correct/one incorrect triplet and two incorrect ones.

2. **Decision Phase.** Now we can use the result of the testing phase to decide whether the given triplet is a D-H triplet or not.
   (a) Suppose the attacker can distinguish between two correct DH triplets and one correct/one random triplet. Then she can input a randomized sequence $[g^{a_0 s}, g^{a_1 t}]$, $[g^{x_0 u}, g^{x_1 v}]$ and $(g^{a_0 a_1 st}, Z^{uv})$ where $[g^{x_0}, g^{x_1}, Z]$ is the target Decision Diffie-Hellman triplet, to the MDH oracle. If the behavior on these inputs is different from the behavior when fed with a sequence of two randomized correct triplets, conclude that the target DDH triplet is *an incorrect Diffie-Hellman triplet.* Else, conclude that it is a correct Diffie-Hellman triplet.
   In other words, the attacker uses the oracle to estimate the following difference:

$$\begin{aligned}\Delta([a,b,1,1],\ [(a,b),(x,y),1,i]) = \\ \mid \Pr[A([g^{a_0}, g^{b_0}],[g^{a_1}, g^{b_1}], g^{a_r b_r}, g^{a_{\bar{r}} b_{\bar{r}}}) = r] - \\ \Pr[A([g^{a_0}, g^{b_0}],[g^{x}, g^{y}], X, Y) = r] \mid ,\end{aligned}$$

   where $X, Y \in_R \{g^{a_0 b_0}, Z\}$ and i is 1 or 0 depending on whether the candidate triplet is a correct or incorrect D-H triplet respectively. We implicitly assume here that the inputs to the oracle are randomized as described above. The estimate of the difference is given with accuracy $\frac{1}{32n^c}$. Now if $Z \neq g^{xy}$ then, as we know from the testing phase, the actual difference is at least $\frac{1}{8n^c}$ and the estimate must be larger than $\frac{1}{8n^c} - \frac{1}{32n^c} = \frac{3}{32n^c}$. Otherwise the actual difference would be 0 and the estimate would be smaller than $\frac{1}{32n^c}$. So depending on the estimate (greater than $\frac{3}{32n^c}$ or smaller than $\frac{1}{32n^c}$) the attacker decides whether the input is an incorrect, or respectively a correct Diffie-Hellman triplet.
   (b) Otherwise, the oracle is able to distinguish between two random triplets and a correct and a random triplet. Then we can feed the MDH oracle with a randomized sequence $[g^{a_0 s}, g^{b_0 t}]$, $[g^{xu}, g^{yv}]$ and g^{wst},Z^{uv} where $[g^x, g^y, Z]$ is the target Decision Diffie-Hellman triplet and where w does not satisfy the equation $w \equiv a_0 b_0 \pmod{Q}$. If the behavior on these inputs is different from the behavior when fed with a sequence of two random triplets, conclude that the target DDH triplet is *a correct Diffie-Hellman triplet.* Else conclude that it is an incorrect Diffie-Hellman triplet.
   In particular, the attacker uses the oracle to estimate the following difference:

$$\begin{aligned}\Delta([(x,y),\ (a,b), i, 0],[a,b,0,0]) = \mid \Pr[A([g^{x}, g^{y}],[g^{a_0}, g^{b_0}], X, Y) = r] \\ - \Pr[A([g^{a_0}, g^{b_0}],[g^{a_1}, g^{b_1}], g^{z_0}, g^{z_1}) = r] \mid ,\end{aligned}$$

where $X, Y \in_R \{Z, g^{z_2}\}, z_0, z_1, z_2 \in_R G_Q$, and i is 1 or 0 depending on whether the candidate triplet is a correct or incorrect D-H triplet respectively. The estimate is given with accuracy $\frac{1}{16n^c}$. Now if $Z = g^{xy}$ then, as we know from the testing phase, the actual difference is at least $\frac{1}{4n^c}$ and the estimate must be larger than $\frac{1}{4n^c} - \frac{1}{16n^c} = \frac{3}{16n^c}$. Otherwise the actual difference would be 0 and the estimate would be smaller than $\frac{1}{16n^c}$, as analyzed in the testing phase above. So depending on the estimate (greater than $\frac{3}{16n^c}$ or smaller than $\frac{1}{16n^c}$) the attacker decides whether the input is a correct, or respectively an incorrect Diffie-Hellman triplet.

4 Universal Malleability Implies Matching = Decision

In this section we will show that for some special classes of cryptosystems indistinguishability of encryptions is equivalent to being able to decide whether a given ciphertext corresponds to a given plaintext. More precisely, we know [Gol93] that indistinguishability of encryptions is equivalent to semantic security. That is, if some information is leaked from the ciphertext then two ciphertext/plaintext pairs can be “matched” (distinguished); and vice-versa. What we do not know, however, is, given that indistinguishability does not hold, *what kind* of information can be extracted about the ciphertext.[3] Here we show that, under certain assumptions about the encryption, if indistinguishability/semantic security does not hold, then given a pair of plaintext and ciphertext it is possible to decide whether the ciphertext comes from this plaintext. Of course this implication only makes sense in either symmetric encryption or probabilistic asymmetric encryption, since in deterministic asymmetric encryption it is straightforward to make this decision: simply encrypt the plaintext and compare to the candidate ciphertext.

We begin by reiterating the definition of indistinguishability of encryptions.

Definition 3. (encryption secure in the sense of indistinguishability) *An encryption scheme (G, E, D) is said to be* secure in the sense of indistinguishability *if, for every probabilistic polynomial time algorithm F (for “Find”), for every probabilistic polynomial time algorithm A, for every constant $c > 0$ and for every sufficiently large n,*

$$Pr\left[F(1^n) = (\alpha, \beta, \gamma) \ s.t. \ \Omega(\alpha, \beta, \gamma) > \frac{1}{n^c}\right] < \frac{1}{n^c} \ ,$$

with

$$\Omega(\alpha, \beta, \gamma) = \left|Pr\left\{A((\gamma), E_{G(1^n)}(\alpha)) = 1\right\} - Pr\left\{A(\gamma, E_{G(1^n)}(\beta)) = 1\right\}\right| \ ,$$

where the probability is taken over the coin tosses of F, A, E and G.

[3] Of course the existing proofs of equivalence between semantic security and indistinguishability [Gol93] constructively extract some information, but this is limited to a specially fabricated function of some specified plaintext/ciphertext pairs.

For our purposes, we need an additional assumption about the encryption scheme. Intuitively, we need to be able to "randomize" any plaintext/ciphertext pair, such that the resulting pair can obtain all possible values. We name the encryption schemes that satisfy this property "universally malleable," to be contrasted to non-malleable schemes [DDN91] that prohibit altering of the ciphertext. The formal definition follows.

Definition 4. (Universal malleability) *An encryption scheme* (G, E, D) *is called* universally malleable *if for all but a negligible part of plaintext-ciphertext pairs* $(a, E_{G(1^n)}(a)) \in (X_n, Y_n)$*, there is a random variable* Z_n *and a p.p.t. TM* T *such that*

- *for every* $z \in Z_n$*, it holds that* $T(a, E_{G(1^n)}(a), z) = (b, E_{G(1^n)}(b))$*, and*
- *for all but a negligible part of pairs* $(c, d) \in (X_n, Y_n)$ *there is a* $z' \in Z_n$ *such that* $T(a, E_{G(1^n)}(a), z') = (c, d)$*.*

Remark: this definition may seem too restrictive, but in fact there are several encryption schemes, at times provably semantically secure under some assumptions, which satisfy it. Examples include the ElGamal cryptosystem [ElG85], the Okamoto-Uchiyama cryptosystem [OU98], the Naccache-Stern higher-order residue cryptosystem [NS98], and the Goldwasser-Micali quadratic-residue cryptosystem [GM84]. Typically this property is derived from the random self reducibility of some underlying problem in which the encryption is based on—be it quadratic or higher order residuosity, the Diffie-Hellman problem, or factoring.

We now proceed to formally define and prove our statement. Again we work in the uniform model of computation.

Theorem 2. *Assume that a* universally malleable *encryption scheme* (G, E, D) *is not secure in the sense of indistinguishability. Then there exists a p.p.t. TM which, given a pair* $(a, E_{G(1^n)}(b))$*, can decide non-negligibly better than random guessing whether* $a = b$*.*

Proof. If an encryption is not secure in the sense of indistinguishability then there exists a p.p.t. adversarial algorithm A (which can be seen as the "oracle" that "breaks" the encryption scheme), a (polynomial) random variable Z_n and two independent (polynomial) random variables (X_n, Y_n) that have the same distribution, such that:

$$\exists\ c > 0, \exists\ N, \text{s.t. for infinitely many } n > N\ , \Pr[(X_n Y_n Z_n) \in B_n^c] > \frac{1}{n^c}\ , \text{where}$$

$$B_n^c = \left\{(\alpha, \beta, \gamma) : \left|\Pr\left[A(\gamma, E_{G(1^n)}(\alpha)) = 1\right] - \Pr\left[A(\gamma, E_{G(1^n)}(\beta)) = 1\right]\right| > \frac{1}{n^c}\right\},$$

where the probabilities are taken over the coin tosses of the key generating algorithm G, the encryption algorithm $E_{G(1^n)}$, the adversarial algorithm A, and the selection of (α, β, γ).

We will show how this adversarial algorithm can be used by an attacker to decide whether $a = b$, in the given pair $(a, E_{G(1^n)}(b))$. For simplicity, we will write $(a, E(b))$.

The process requires three phases (including the preparation phase).

1. **Preparation phase.** In this phase the attacker finds two plaintexts whose ciphertexts she can distinguish. This is possible given the assumptions on the power of the adversarial algorithm A above.
Specifically, the attacker chooses a random message pair (m_0, m_1) from the distribution X_n and tries to estimate the following probability:
$$\Pr[A([m_0, m_1], [E(m_r), E(m_{\bar{r}})]) = r] - \frac{1}{2}$$
where $r \in_R \{0, 1\}, r \oplus \bar{r} = 1$
with accuracy better than $\frac{1}{2n^c}$. Now if the estimate is greater than $\frac{3}{2n^c}$, the actual probability is greater than $\frac{3}{2n^c} - \frac{1}{2n^c} = \frac{1}{n^c}$ and the message pair is selected for the next step. Otherwise it is rejected and a new pair is selected. The number of experiments needed to estimate this probability with accuracy $\frac{1}{2n^c}$ is polynomially bounded since the encryption scheme is not secure in the sense of indistinguishability, and it can be computed using the Hoefding inequality.
Note that the estimation is performed by randomizing the input to algorithm A. This is where the property of *universal malleability* is crucial, in guaranteeing that the randomization will always succeed and that the randomized input can take all possible values: recall from the definition that for every $(m, E(m)) \in (X_n, Y_n)$ and $z' \in_R Z_n$ it holds that $T(m, E(m), z') = (b, E(b))$ for some $b \in X_n$, for all but a negligible part of plaintexts m. Therefore we can randomize the input sequence to the oracle. Now, from the second part of universal malleability we have that for all but a negligible part of pairs $(c, d) \in (X_n, Y_n)$ there is a $z' \in Z_n$ such that $T(m, E(m), z') = (c, d)$. Thus a randomization of $(m, E(m))$ achieved by T choosing a random $z' \in_R Z_n$ as above, results in a uniformly chosen pair from the distribution (X_n, Y_n), and all but a negligible fraction of those pairs can be generated by T in this manner.
Let m_0 and m_1 be the two messages that the algorithm can distinguish. We denote this as follows:
$$\Pr[A([m_0, m_1], [E(m_r), E(m_{\bar{r}})]) = r] - \frac{1}{2} > \frac{1}{n^c} , \qquad (3)$$
where $r \in_R \{0, 1\}, r \oplus \bar{r} = 1$.
2. **Testing phase.** As in section 3, assume that the attacker uniformly and independently selects $r \in_R \{0, 1\}, m_2, m_3 \in_R X_n, v, w \in_R Y_n$, and then uses the oracle to estimate:
$$\left| \Pr[A([m_2, m_3], [v, w]) = r] - \frac{1}{2} \right| ,$$

where v, w do not encrypt m_2 nor m_3. Then again it is clear that the probability of the oracle in finding the attacker's random coin tosses for r is negligible (we formalize it as less than $\frac{1}{2n^c}$), since no information (in the Shannon sense) about r is included in the input of A. Thus, combining equation (3), we have the equivalent of lemma 1:

Lemma 2. *For every $m_2, m_3 \in_R X_n, r \in_R \{0,1\}, r \oplus \bar{r} = 1, v, w \in_R Y_n$, and for large enough n it holds that*

$$\begin{aligned}\Delta([m_0, m_1, 1, 1], [m_2, m_3, 0, 0]) = |\ &Pr\,[A([m_0, m_1], [E(m_r), E(m_{\bar{r}})]) = r] - \\ &Pr\,[A([m_2, m_3], [v, w]) = r]| > \frac{1}{2n^c}\end{aligned}$$

Now the attacker runs algorithm A in order to estimate the difference:

$$\begin{aligned}\Delta([m_0, m_1, 1, 1], [m_4, m_5, 1, 0]) = |\ &\Pr\,[A([m_0, m_1], [E(m_r), E(m_{\bar{r}})]) = r] - \\ &\Pr\,[A([m_4, m_5], [X, Y]) = r]|\ ,\end{aligned}$$

where $X, Y \in_R \{E(m_4), t\}, t \in_R Y_n, r \in_R \{0,1\}, \bar{r} \oplus r = 1$.
The estimate is given with accuracy $\frac{1}{16n^c}$; as in the preparation phase, we use the property of universal malleability to allow A to randomize the input and run the oracle as many times as dictated by the Hoefding inequality. As in section 3 we can test if the difference here is significant, i.e., greater than or equal to $\frac{3}{16n^c}$. In this case, the actual difference is at least $\frac{3}{16n^c} - \frac{1}{16n^c} = \frac{1}{8n^c}$, i.e. the attacker is able to distinguish between two correct plaintext/ciphertext pairs, and one correct and one incorrect one. If on the other hand the estimate is smaller than $\frac{3}{16n^c}$, the actual difference will be smaller than $\frac{3}{16n^c} + \frac{1}{16n^c} = \frac{1}{4n^c}$, and as in section 3 we can show that $\Delta([m_4, m_5, 1, 0], [m_2, m_3, 0, 0])$ will be greater than $\frac{1}{4n^c}$. In other words, the difference between two incorrect plaintext/ciphertext pairs, and one incorrect and one correct pair has to be significant. This is shown as follows:

$$\begin{aligned}\frac{1}{2n^c} &< \Delta([m_0, m_1, 1, 1], [m_2, m_3, 0, 0]) \\ &= |\Pr[A([m_0, m_1], [E(m_r), E(m_{\bar{r}})]) = r] - \Pr[A([m_2, m_3], [v, w]) = r]| \\ &= |\Pr[A([m_0, m_1], [E(m_r), E(m_{\bar{r}})]) = r] - \Pr[A([m_4, m_5], [X, Y]) = r] \\ &\quad + \Pr[A([m_4, m_5], [X, Y]) = r] - \Pr[A([m_2, m_3], [v, w]) = r]| \\ &\le |\Pr[A([m_0, m_1], [E(m_r), E(m_{\bar{r}})]) = r] - \Pr[A([m_4, m_5], [X, Y]) = r]| \\ &\quad + |\Pr[A([m_4, m_5], [X, Y]) = r] - \Pr[A([m_2, m_3], [v, w]) = r]| \\ &= \Delta([m_0, m_1, 1, 1], [m_4, m_5, 1, 0]) + \Delta([m_4, m_5, 1, 0], [m_2, m_3, 0, 0])\end{aligned}$$

Thus, for uniformly chosen $m_4, m_5 \in_R X_n$, i.e., $\Pr[(m_4, m_5)] = \frac{1}{|X_n|^2}$, and for j enumerating all possible pairs, we have:

$$\Sigma_j\ [\Delta([m_0, m_1, 1, 1], [m_4, m_5, 1, 0]) +$$

$$\begin{aligned}
\Delta([m_4, m_5, 1, 0], [m_2, m_3, 0, 0])] &> \Sigma_j \frac{1}{2n^c} \iff \\
\Sigma_j\ \Delta([m_0, m_1, 1, 1], [m_4, m_5, 1, 0]) &+ \\
\Sigma_j\ \Delta([m_4, m_5, 1, 0], [m_2, m_3, 0, 0]) &> |X_n|^2 \frac{1}{2n^c} \iff \\
\Sigma_j\ \frac{\Pr[(m_4, m_5)]}{\Pr[(m_4, m_5)]} \Delta([m_0, m_1, 1, 1], [m_4, m_5, 1, 0]) &+ \\
\Sigma_j\ \frac{\Pr[(m_4, m_5)]}{\Pr[(m_4, m_5)]} \Delta([m_4, m_5, 1, 0], [m_2, m_3, 0, 0]) &> |X_n|^2 \frac{1}{2n^c} \iff \\
|X_n|^2 \Sigma_j\ \Pr[(m_4, m_5)] \Delta([m_0, m_1, 1, 1], [m_4, m_5, 1, 0]) &+ \\
|X_n|^2 \Sigma_j\ \Pr[(m_4, m_5)] \Delta([m_4, m_5, 1, 0], [m_2, m_3, 0, 0]) &> |X_n|^2 \frac{1}{2n^c} \iff \\
\Sigma_j\ \Pr[(m_4, m_5)] \Delta([m_0, m_1, 1, 1], [m_4, m_5, 1, 0]) &+ \\
\Sigma_j\ \Pr[(m_4, m_5)] \Delta([m_4, m_5, 1, 0], [m_2, m_3, 0, 0]) &> \frac{1}{2n^c} \iff \\
E[\Delta([m_0, m_1, 1, 1], [m_4, m_5, 1, 0])] &+ \\
E[\Delta([m_4, m_5, 1, 0], [m_2, m_3, 0, 0])] &> \frac{1}{2n^c}\ ,
\end{aligned}$$

where the expected values are taken over the choice of the pair (m_4, m_5). Therefore, if $E[\Delta([m_0, m_1, 1, 1], [m_4, m_5, 1, 0])] < \frac{1}{4n^c}$ then it must be that $E[\Delta([m_4, m_5, 1, 0], [m_2, m_3, 0, 0])] > \frac{1}{4n^c}$.

This concludes the proof that the oracle may either distinguish between two correct plaintext/ciphertext pairs and one correct/one incorrect pair, *or* between one correct/one incorrect pair and two incorrect plaintext/ciphertext pairs.

We now proceed to the last phase.

3. **Decision phase.** In this phase we use the result of the testing phase accordingly. If the first difference is significant, then the attacker estimates the difference

$$\begin{aligned}
\Delta([m_0, m_1, 1, 1], [m_4, a, 1, i]) = |\ &\Pr\left[A([m_0, m_1], [E(m_r), E(m_{\bar{r}})]) = r\right] - \\
&\Pr\left[A([m_4, a], [X, Y]) = r\right]|\ ,
\end{aligned}$$

where $X, Y \in_R \{E(m_4), E(b)\}$. The estimate is given with accuracy $\frac{1}{32n^c}$, after polynomially many trials. In each trial each plaintext/ciphertext pair is randomized over all (but a negligible part) of possible combinations $(c, d) \in (X_n, Y_n)$; again we utilize the universal malleability property to do this, as in the preparation phase. Now as we know from the testing phase if $a \neq b$ the actual difference is at least $\frac{1}{8n^c}$, so the estimate in this case must be greater or equal to $\frac{1}{8n^c} - \frac{1}{32n^c} = \frac{3}{32n^c}$. Otherwise, if $a = b$ the actual difference is 0 and the estimate must be less or equal to $\frac{1}{32n^c}$. Therefore depending on the estimate (greater than $\frac{3}{32n^c}$ or smaller than $\frac{1}{32n^c}$) the attacker decides whether the input is a correct $(i = 1)$ or an incorrect $(i = 0)$ plaintext/ciphertext pair.

Accordingly, if the testing phase showed that the second difference is significant, then the attacker estimates

$$\Delta([a, m_5, i, 0], [m_2, m_3, 0, 0]) = |\ \Pr\left[A([a, m_5], [X, Y]) = r\right] - \\ \Pr\left[A([m_2, m_3], [v, w]) = r\right]|\ ,$$

where $X, Y \in_R \{E(b), t\}, t, v, w \in_R Y_n$. Here the required accuracy is $\frac{1}{16n^c}$. If $a = b$ then from the testing phase we know that the actual difference is at least $\frac{1}{4n^c}$, and therefore the estimate will be larger than $\frac{1}{4n^c} - \frac{1}{16n^c} = \frac{3}{16n^c}$. Otherwise, when $a \neq b$, the difference would be 0 and the estimate would be smaller than $\frac{1}{16n^c}$. Here again, the attacker can decide whether the input to algorithm A is a correct $(i = 1)$ or an incorrect $(i = 0)$ plaintext/ciphertext pair depending on the value of the estimated difference (above $\frac{3}{16n^c}$ or below $\frac{1}{16n^c}$). Again note that *universal malleability* is fundamental to this proof in order to be able to feed the oracle with a randomized input sequence.

5 Extensions

The original version of the matching Diffie-Hellman problem, defined in [FTY96], was slightly different from the one used in the analysis above. For convenience we name it "matching D-H II".

Definition 5. (Matching Diffie-Hellman Problem II) *For security parameter n, for uniformly chosen $a_i, b_i \in_R G_Q$ $(i \in \{0, 1\})$, P a prime with $|P - 1| = \delta + n$ for a specified constant δ, and for $g \in Z_P^*$ a generator of prime order $Q = (P - 1)/\gamma$ for a specified small integer γ, given $[g^{a_0}, g^{a_0 b_0}]$, $[g^{a_1}, g^{a_1 b_1}]$ and g^{b_r},$g^{b_{\bar{r}}}$, $r, \bar{r} \in_R \{1, 0\}, r \oplus \bar{r} = 1$, find r with probability better than $\frac{1}{2} + \frac{1}{n^c}$ for some constant c for large enough n.*

Using the same techniques of section 3 it can be shown that this version of the problem is also equivalent to the decision Diffie-Hellman problem, and therefore the two versions of the "matching" problem are equivalent.

References

[Can97] R. Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In B. Kaliski, editor, *Advances in Cryptology — CRYPTO '97 Proceedings, LLNCS 1294*, pages 455–469, Santa Barbara, CA, August 17–21 1997. Springer-Verlag.

[CFT98] A. Chan, Y. Frankel, and Y. Tsiounis. Easy come–easy go divisible cash. In *Advances in Cryptology — Proceedings of Eurocrypt '98 (Lecture Notes in Computer Science 1403)*, pages 561–575, Helsinki, Finland, May 31–June 4 1998. Springer-Verlag. International patent pending. Available at http://www.ccs.neu.edu/home/yiannis/pubs.html.

[CS98] R. Cramer and V. Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *Advances in Cryptology: Crypto '98, Proceedings (Lecture Notes in Computer Science 1462)*, pages 13–25, 1998. Available at http://www.cs.wisc.edu/ shoup/papers/.

[Dam91] I. B. Damgård. Towards practical public key systems against chosen ciphertext attacks. In J. Feigenbaum, editor, *Advances in Cryptology, Proc. of Crypto '91 (Lecture Notes in Computer Science 576)*, pages 445–456. Springer-Verlag, 1991.

[DDN91] O. Dolev, C. Dwork, and M. Naor. Non-malleable cryptography. In *Proceedings of the 23rd Annual Symposium on Theory of Computing, ACM STOC*, 1991.

[ElG85] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. Inform. Theory*, 31:469–472, 1985.

[FTY96] Y. Frankel, Y. Tsiounis, and M. Yung. Indirect discourse proofs: achieving fair off-line e-cash. In *Advances in Cryptology, Proc. of Asiacrypt '96 (Lecture Notes in Computer Science 1163)*, pages 286–300, Kyongju, South Korea, November 3–7 1996. Springer-Verlag. International patent pending. Available at http://www.ccs.neu.edu/home/yiannis/pubs.html.

[FTY98] Y. Frankel, Y. Tsiounis, and M. Yung. Fair off-line cash made easy. In *Advances in Cryptology, Proc. of Asiacrypt '98 (Lecture Notes in Computer Science)*. Springer-Verlag, October 18–22 1998. To appear. Available at http://www.ccs.neu.edu/home/yiannis/pubs.html.

[GM84] S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, April 1984.

[Gol93] O. Goldreich. A uniform-complexity treatment of encryption and zero-knowledge. *Journal of Cryptology*, 6(1):21–53, 1993. Available at http://www.wisdom.weizmann.ac.il/people/homepages/oded.

[NS98] D. Naccache and J. Stern. A new cryptosystem based on higher residues. In *ACM CCS '98—Communications and Computer Security*, 1998. To appear.

[Oka95] T. Okamoto. An efficient divisible electronic cash scheme. In Don Coppersmith, editor, *Advances in Cryptology, Proc. of Crypto '95 (Lecture Notes in Computer Science 963)*, pages 438–451. Springer-Verlag, 1995. Santa Barbara, California, U.S.A., August 27–31.

[OU98] T. Okamoto and S. Uchiyama. An efficient public-key cryptosystem. In *Advances in Cryptology – Eurocrypt 98 proceedings (Lecture Notes in Computer Science 1403)*, pages 308–318, Espoo, Finland, May 31–June 4 1998. Springer-Verlag.

[TY98] Y. Tsiounis and M. Yung. On the security of El Gamal-based encryption. In *International workshop on Public Key Cryptography (PKC '98) (Lecture Notes in Computer Science 1431)*, pages 117–134, Yokohama, Japan, February 5-6 1998. Springer-Verlag. Available at http://yiannis.home.ml.org.