

Empirical study: What are the factors that influence the security of the Diffie-Hellman exchange?

No Author Given

No Institute Given

Abstract. The Diffie-Hellman public key exchange protocol is a gold mine in cryptography. It is so fundamental and yet so hard to study that countless papers concerned with its security are available in literature. Most of these papers treat the underlying group as any generic group and try to draw general conclusions, as the size of the group goes to infinity. This paper is concerned with practical aspects of the exchange. We recognize that in practice we cannot work with infinite size groups and therefore the structure of the underlying group must be crucial to the exchange. We give compelling empirical evidence of this fact. The groups under study are small since we need to study the structure of the exchange in totality – for all inputs, yet the conclusions we draw seem to scale for larger group sizes, in the end the hope is to relate the security of the exchange with only a few defining elements of the underlying group (isomorphic structures).

Keywords and phrases: public key cryptography, statistical indistinguishability, group theory, prime subgroups, group structure.

1 Introduction.

In this article we try to relate the structure of the underlying group with the security of the Diffie-Hellman public key exchange protocol.

In its most basic form, the Diffie-Hellman protocol chooses a finite cyclic group (G, \cdot) of order N , with generator g , where \cdot denotes the group operation. The group G is $G = \{g^0, g^1, \dots, g^{N-1}\}$ or symbolically $G = \langle g \rangle$. Note that G , g and N are public information.

The participants in the information transfer \mathcal{A} and \mathcal{B} each randomly and independently chooses an integer $a \in [1, N]$ and $b \in [1, N]$ independently. Then \mathcal{A} computes g^a , \mathcal{B} computes g^b and exchange these elements of G over an insecure channel. Since each of \mathcal{A} and \mathcal{B} knows their respective a and b they can compute g^{ab} , which or a publicly known derivation $K_{\mathcal{A}, \mathcal{B}}$ of that becomes the public key. Any method of converting g^{ab} to $K_{\mathcal{A}, \mathcal{B}}$ is publicly known, and the security of the key $K_{\mathcal{A}, \mathcal{B}}$ is directly dependent on the security of g^{ab} , therefore for the sake of simplicity we will consider g^{ab} as the established key of the exchange for the rest of this paper.

The security of the exchange is concerned with the ability of an outside observer (reading the exchanged numbers g^a and g^b) to deduce information about the key g^{ab} .

Immediately after its discovery, studies about the protocol's security have concentrated in analyzing the problem using two classical formulations (the discrete logarithm and the computational Diffie-Hellman problems):

Assumption 1 (DL) *For a cyclic group G , generated by g , we are given g and g^n , $n \in \mathbb{N}$, then it is hard to calculate n .*

Assumption 2 (CDH) *Given g, g^a, g^b it is hard to compute g^{ab} .*

Clearly, if these assumptions are not satisfied then \mathcal{C} , an adversary¹, can gain access to the key g^{ab} . The relationship between these two assumptions has been extensively studied. It is clear that the CDH assumption will not be satisfied in a group where finding the discrete logarithm solution is easy. In [10], [2], the authors show that in several settings the validity of the CDH assumption and the hardness of the discrete logarithm problem are in fact equivalent.

Unfortunately, the DL and the CDH assumptions are not enough to ensure security of the Diffie-Hellman key exchange protocol. Even if these assumptions are true, the eavesdropper \mathcal{C} may still be able to gain useful information about g^{ab} . For example, if \mathcal{C} can predict 90% of the bits in g^{ab} with high probability then for all intents and purposes the key exchange protocol is broken. Moreover, there exist protocols where the knowledge of even one bit will break its security (e.g., Casino electronic games). With the current state of knowledge we cannot be confident that assuming only CDH, a scenario like the one described above does not exist ([1]).

For this reason, subsequent studies formulate and are study the validity of an assumption of the following type (decision Diffie-Hellman problem):

Assumption 3 (DDH) *Given g, g^a, g^b and an element $z \in G$ it is hard to decide whether or not $z = g^{ab}$.*

In this form the DDH assumption constitutes a necessary condition for the security of the Diffie-Hellman key exchange protocol. Furthermore, in [8] the authors construct groups based on elliptic curves where the DDH assumption is not satisfied while the CDH and the discrete logarithm problem are proven to be equivalent and hard. This fact prompts the necessity to directly check the validity of the DDH assumption for a given group.

The formulation above is not specific enough to warrant a specific method for checking it. Therefore, a more specific assumption is needed. Indeed, this fact materializes itself in many recent papers, we shall name only a few: [3, 4, 6, 11], which call this new form of the assumption the Diffie Hellman Indistinguishability assumption (DHI). We note that [7, 8] use the same form except it continues to call it DDH. We will phrase this assumption as follows:

Assumption 4 (DHI) *Given g, g^a, g^b the distribution of g^{ab} is indistinguishable from the Discrete Uniform distribution on the elements of G ($DU(G)$).*

The notion of indistinguishability we use here is the usual statistical one (in contrast to the computational notion previously used). Two variables are indistinguishable if they have almost surely the same distribution, and in our specific context, having a finite set of outcomes (the elements of G) translates into the two variables having exactly the same distribution.

The present article is structured as follows. Section 2 presents the test used to asses the security of the exchange. Section 3 presents numerical results obtained when using the test and statistics that relate the relevant factors of the security of the D-H exchange. Finally, section 4 concludes.

¹ There are various concepts of adversary in cryptographic literature considering the power and authority they have. In this article we assume that our adversary is a passive eavesdropper.

2 Obtaining the test value used to describe the security of the Diffie-Hellman exchange.

In a recent article [5] the authors propose the use of a statistical test that given a group G and a sample of past exchanges will output a measure (based on the Kullback-Leibler divergence measure [9]) of departure from the uniform distribution².

For the sake of completeness we summarize the testing procedure in a few paragraphs, the reader interested in more details and proofs of the results is referred to the cited article.

The hypotheses of the test are:

$$\begin{cases} H_0 : & \text{The distribution of } g^{ab} | (g^a, g^b) \text{ is } DU(G) \\ H_a : & \text{The distribution of } g^{ab} | (g^a, g^b) \text{ is NOT } DU(G) \end{cases} \quad (1)$$

The symbol $X|Y$ means that we refer to the distribution of the variable X conditioned by the σ -algebra generated by the variable Y .

Denote the elements of G as $\{g_1, g_2, \dots, g_N\}$. Suppose we look at the joint distribution of the vector (g^a, g^b, g^{ab}) . Suppose we can look at all the possible triples (g^a, g^b, g^{ab}) when $a, b \in \{1, 2, \dots, N\}$ take all the possible values. Clearly, there are N^2 such possible triples and assuming that a and b are chosen at random, each such triple will have probability $1/N^2$. The last element in the triple g^{ab} will get mapped into N possible values (the elements of G). Thus, some values in G will be repeated. For an element $g_k \in G$ denote m_k the number of times g_k appears in the place of g^{ab} among all the N^2 triples. We have then $\sum_k m_k = N^2$. For any pair (g^a, g^b) that corresponds to $g^{ab} = g_k$ we can then calculate the following conditional probabilities:

$$p(g^a = g_i, g^b = g_j | g^{ab} = g_k) = p(g_i, g_j | g_k) = \frac{1}{m_k} \mathbf{1}_A(g_i, g_j, g_k),$$

where A is the set of all possible N^2 tuples (g^a, g^b, g^{ab}) ,³

For 3 variables X, Y and Z we define the *conditional measure of uncertainty* (Kullback-Leibler divergence measure) as:

$$H(X, Y|Z) = - \sum_{i=1}^n \sum_{j=1}^m \sum_{k=1}^l p(x_i, y_j, z_k) \log p(x_i, y_j | z_k), \quad (2)$$

with the convention $0(-\infty) = 0$. In our case considering the variables $X = g^a, Y = g^b$ and $Z = g^{ab}$ we obtain:

² A case could be made that the particular divergence measure used is irrelevant, any statistical procedure measuring departure from the uniform distribution would work just as well.

³ We use the notation $\mathbf{1}_A(x)$ to denote the indicator function of the set $A \subset \Omega$, i.e., $\mathbf{1}_A : \Omega \rightarrow \{0, 1\}$ is given by:

$$\mathbf{1}_A(x) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{if } x \notin A \end{cases}$$

$$\begin{aligned}
H(g^a, g^b | g^{ab}) &= - \sum_{i=1}^N \sum_{j=1}^N \sum_{k=1}^N p(g_i, g_j, g_k) \log p(g_i, g_j | g_k) \\
&= - \sum_{k=1}^N \sum_{i,j=1}^N \frac{1}{N^2} \log \frac{1}{m_k} \mathbf{1}_A(g_i, g_j, g_k) \\
&= - \sum_{k=1}^N \frac{m_k}{N^2} \log \frac{1}{m_k}. \tag{3}
\end{aligned}$$

Under the null hypothesis H_0 , the distribution of $(g^{ab} | g^a, g^b)$ is uniform, therefore the m_k multiplicities are equal if H_0 is true. This automatically implies that $m_k = N$ for all k 's and then the entropy function in (3) is:

$$H(g^a, g^b | g^{ab}) = - \sum_{k=1}^N \frac{1}{N} \log \frac{1}{N} = \log N$$

The testing statistics is thus:

$$T_N = H(g^a, g^b | g^{ab}) - \log N = \sum_{k=1}^N \frac{m_k}{N^2} \log m_k - \log N. \tag{4}$$

This test is based on the whole set of values in G^2 . Accordingly, if the value of the test equals zero then the null hypothesis H_0 is true, any other value of the test will support the alternative hypothesis. We summarize this result in the following:

Lemma 1 (Testing Procedure). *With the previous notations if $T_N = 0$ then the DHI assumption is satisfied in a given group G .*

We exemplify the use of the test procedure in the following section.

3 Numerical results of the test.

In this section we look at small cyclic groups of the type (\mathbb{Z}_p^*, \cdot) , where \cdot denotes multiplication modulo p . We first look to (\mathbb{Z}_p^*, \cdot) themselves, for p primes in two separate ranges $p \in (2000, 4000)$ and $p \in (9000, 11000)$.

We calculate the test values for each such group (for such small groups we do not need to estimate or construct statistical distributions for the test values), with the idea to compare the groups themselves from the perspective of the test and identify (if possible) patterns. To our knowledge this is the first approach of this kind.

It is known – due to the existence of the Pohlig-Hellman algorithm⁴ that in all of these groups the Discrete Logarithm problem is easy and therefore the Diffie Hellman exchange should be breakable. It is also *conjectured* that the actual security depends on the size of the largest factor in the

⁴ which computes the Legendre symbol in these groups and therefore gives a distinguisher against DDH (see [7])

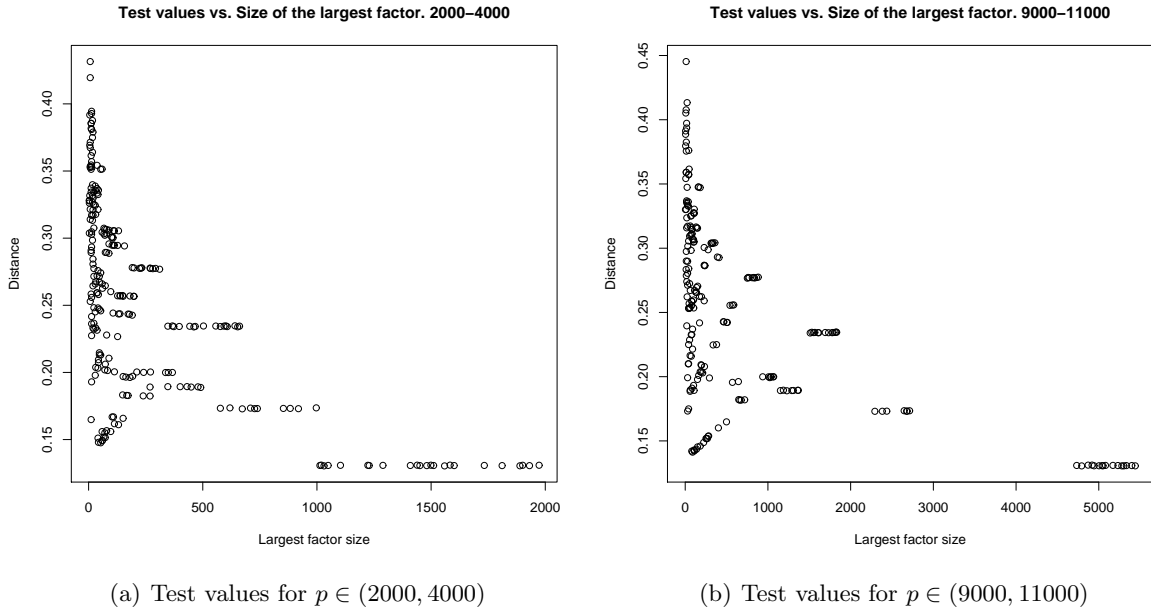


Fig. 1. On the x axis we plot the biggest factor in the decomposition of $p - 1$. Values closer to zero on the y -axis represent safer groups for DH exchange.

decomposition of $p - 1$ ⁵. For this reason it is believed that the “most secure” groups among (\mathbb{Z}_p^*, \cdot) are the multiplicative groups \mathbb{Z}_p^* generated by the so called safe primes (i.e., primes p with $p = 2q + 1$ and q another prime). We shall call any such group a *safe group*.

We expect that these facts will translate to our test as well.

Figure 1 on page 5 presents the test values vs. the size of the largest prime factor in the decomposition of $p - 1 = q_1 q_2 \dots q_n$. We can immediately see that the structure of the test values for the two ranges is very similar.

In both of these images, points corresponding to values closer to 0 on the y axis represent groups that are more secure for the DH exchange.

Note that while the points in the lower right corner of the image correspond indeed to the safe primes and they are clearly more secure than the other groups as the popular belief would tell us, we can also see that there exist certain groups which have a small factor (lower left corner) and yet they are comparably secure.

This fact is investigated in the Figure 2 on page 6 where we plot the test values obtain for a group \mathbb{Z}_p versus the number of factors in the decomposition of $p - 1$. While we can see more clearly now that the groups corresponding to the safe primes ($x = 2$ factors in the plot) are indeed more secure than all the other groups, we also find that generally as the number of factors in the decomposition increases the security decreases.

Based on these two plots it would seem that both the number of factors and the size of the largest factor are important elements when considering the security of the exchange.

⁵ This is due to the nature of the algorithm

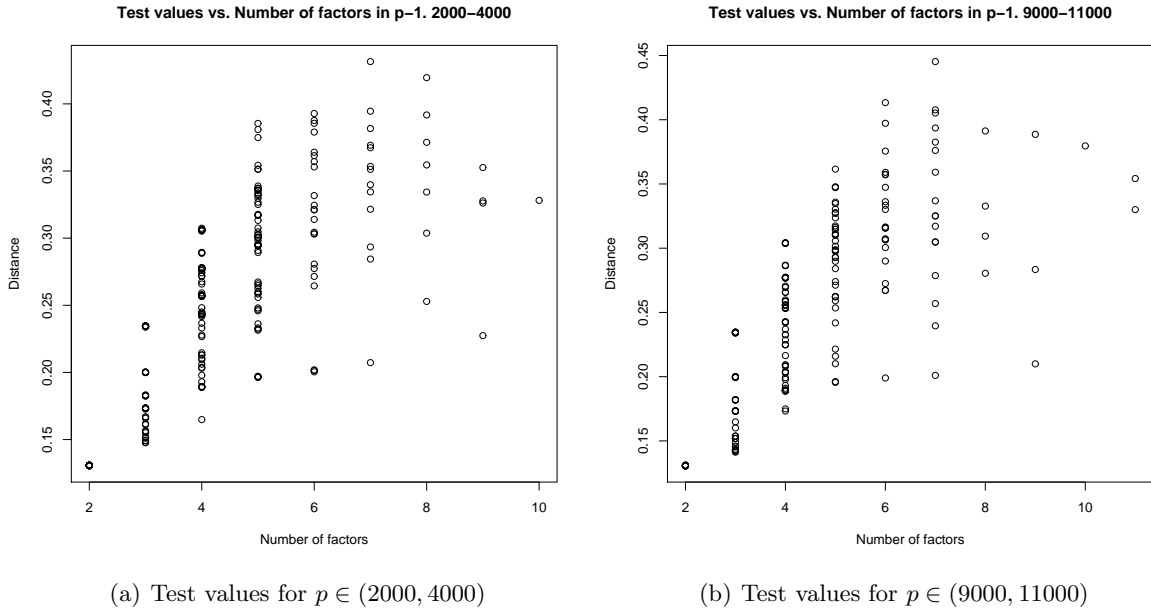


Fig. 2. On the x axis we plot the number of factors in the decomposition of $p - 1$. Values closer to zero represent safer groups for DH exchange.

But this now becomes a statistical problem: trying to relate two determining factors to the variable that quantifies the security of the exchange. There probably exist other factors that are important but let us concentrate on these two for the current paper. We know from the statistical theory that if there would be no interaction between the number of factors in the decomposition and the size of the largest factor then we should see points inside each category close to parallel lines. For exemplification we plotted in Figure 3 on page 7 the same image as in Figure 1, but with the points separated by the number of factors in each group. We eliminated the safe groups from the comparison and we only made the picture for $p \in (9000, 11000)$ since for the other range the image looks very similar.

We can start to see that there must be interaction between the two factors. To exemplify better we separated the points depending on the number of factors and we plotted them in Figure 4. We see better that the determining elements for the security of the DH exchange seem to be correlated (they are interacting).

As an example of such discrepancy the group \mathbb{Z}_{9473} which has order $9473 - 1 = 2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 37$, and whose biggest factor in the decomposition is 37 is more secure (test value 0.2) than both groups: \mathbb{Z}_{9421} and \mathbb{Z}_{9781} , whose decompositions of $p - 1$ are $9421 - 1 = 2 \times 2 \times 3 \times 5 \times 157$ and $9781 - 1 = 2 \times 2 \times 3 \times 5 \times 163$ respectively⁶.

Next we have analyzed statistically the relationship between the test values that quantify the strength of the relationship and the size of the largest factor in the decomposition of $p - 1$ (treated as a quantitative variable) and the number of factors in the same decomposition (treated as a categorical variable). We included interaction terms in the model and we present the ANOVA table

⁶ the test values obtained for these two later groups are very close to each other 0.3475897 and 0.3476914.

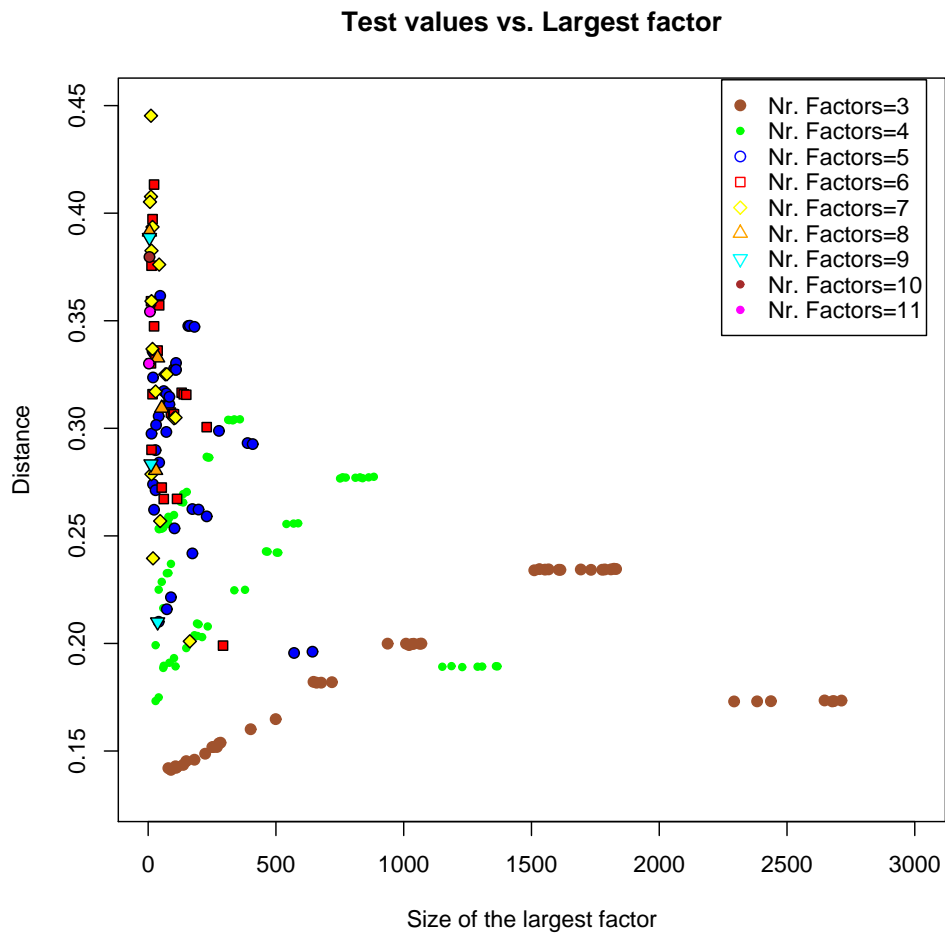


Fig. 3. This is the same image as Figure 1(b) but with points corresponding to number of factors in the decomposition of $p - 1$ identified.

(Table 1) on page 7. Afterwards, in Table 2 we present the estimated coefficients of the regression lines for each level, and for each, a test of whether the mean is actually zero. There are 218 primes between 9,000 and 11,000.

Table 1: ANOVA table for the relationship between the size of the largest factor, the number of factors and the test values

Factors	df	Deviance	df	Residual	Deviance
Largest factor	1	0.37997	216		0.77629
Number of Factors	9	0.41977	207		0.35652
Interaction	8	0.08881	199		0.26772
Continued on next page					

Table 1 – Continued from previous page			
Factors	df	Deviance	df Residual Deviance
Error		217	1.15626

Table 2: Effects for each level of factor. The semicolon denotes the levels of the interaction term. The individual factors have to be included even though they appear not significant since the interaction is.

Factor levels	Estimate	Std. Error	t-values	p-values
(Intercept)	1.322e-01	2.221e-01	0.595	0.552573
Largest factor	-2.718e-07	4.353e-05	-0.006	0.995025
Nr factors=3	2.794e-02	2.223e-01	0.126	0.900116
Nr factors=4	1.098e-01	2.222e-01	0.494	0.621938
Nr factors=5	1.766e-01	2.223e-01	0.795	0.427804
Nr factors=6	2.202e-01	2.224e-01	0.990	0.323292
Nr factors=7	2.406e-01	2.225e-01	1.082	0.280778
Nr factors=8	2.462e-01	2.253e-01	1.093	0.275831
Nr factors=9	2.459e-01	2.248e-01	1.094	0.275293
Nr factors=10	2.475e-01	2.249e-01	1.100	0.272574
Nr factors=11	1.798e-01	2.329e-01	0.772	0.440964
maxFact:N.fact=3	2.375e-05	4.399e-05	0.540	0.589972
maxFact:N.fact=4	-7.348e-06	4.505e-05	-0.163	0.870586
maxFact:N.fact=5	-1.260e-04	5.954e-05	-2.115	0.035637 *
maxFact:N.fact=6	-4.137e-04	1.140e-04	-3.629	0.000362 ***
maxFact:N.fact=7	-9.027e-04	2.112e-04	-4.274	2.98e-05***
maxFact:N.fact=8	-1.609e-03	1.060e-03	-1.518	0.130544
maxFact:N.fact=9	-4.757e-03	1.525e-03	-3.119	0.002087 **
maxFact:N.fact=10	NA	NA	NA	NA
maxFact:N.fact=11	6.041e-03	1.297e-02	0.466	0.641823
Signif. codes: '***' = 0.001; '**' = 0.01; '*' = 0.05; '.' = 0.1				

We included interaction terms in the model and we present the ANOVA table (Table 1) on page 7. Afterwards, in Table 2 we present the estimated coefficients of the regression lines for each level, and for each, a test of whether the mean is actually zero. There are 218 primes between 9,000 and 11,000. We note that there was only one prime within the range whose $p - 1$ decomposition had 10 factors thus the interaction for that level could not be estimated.

We can see very clearly from the table that the interaction between the two factors analyzed is significant. We do not present the results for the other range of primes studied 2000 – 4000 since they are entirely similar.

So what is the conclusion to be drawn from these numbers?

These numbers show that the interaction between the number of factors in the decomposition of $p-1$ and the size of the largest factor in the decomposition is statistically significant for the security of the Diffie-Hellman security as quantified by our test.

In plain terms, it would seem natural that as the size of the largest factor in the decomposition increases the group becomes more complex and therefore it is more secure. Likewise, as the number of factors in the decomposition increases, there are more equations to solve modulo each factor therefore having a larger number intuitively would also increase the security.

However, as the results in the table show that is not necessarily so, and since the interaction between the two is significant the combination of the two factors is important and the seemingly logical statements presented are not necessarily true.

3.1 Groups where the exchange is secure.

As we already have said, in all of these groups there exists a distinguisher for the DL problem given by the Pohlig-Hellman algorithm. This is also translated in our approach to the problem by the fact that none of the test values are zero (or non-significantly away from zero).

Perhaps then, these groups are not interesting from the cryptographic perspective. However, in a prime subgroup⁷ of \mathbb{Z}_p no distinguisher is known and it is conjectured that these groups are good for the cryptographic exchange.

We did find evidence that these groups are significantly safer for the DH exchange than any other ones considered. Figure 5 on page 12 presents the results obtained when comparing the prime subgroups of safe groups between \mathbb{Z}_{9000} and \mathbb{Z}_{11000} with the safe groups themselves. To be more specific we took all \mathbb{Z}_p with $p = 2q + 1$, $p \in (9000, 11000)$ (p and q both primes) and constructed one prime subgroup of order q from elements of each such group. Then we calculated the test value for all the subgroups and groups and plotted their histograms in Figure 5 (upper and respectively lower histogram). To make sure that we compare similar numbers, we have estimated rather than calculated the test values using sample sizes of 8×10^7 for each group or subgroup.

The lower histogram in the Figure 5 represents test values for the groups which were designated as the most secure by the previous step. Accordingly, conducting a significance test for the equality of means of two population finds that there is indeed a significant difference (for any significance level α) and that the prime subgroups are indeed more secure than anything else of the type (\mathbb{Z}_p^*, \cdot) .

Furthermore, it would seem from the Figure 5 that the actual size of the group (or subgroup) does not make a difference. This is in fact not true, it just seems that way due to the closeness of the test values obtained. We are exemplifying this issue with the values obtained for the subgroups in Figure 6 on page 13.

We can also see that as the size of the group increases the security of the exchange based on the underlying group increases as well. This shows that if the structure of the underlying group remains the same (in other words all things constant), then the popular assertion that increasing the group size increases the security is indeed true.



⁷ A prime subgroup of a group G of order N is defined as a set of elements $H \subseteq G$ such that $H = \langle g_1 \rangle$ (is generated by the element g_1) and the order of g_1 (denoted with $|g_1|$) is a prime divisor of the order of the group (i.e., $|g_1|$ divides N but $|g_1|^2$ does not divide N).

4 Conclusion.

This paper **does not** break or gives an algorithm to break the Diffie-Hellman exchange. What we do is analyze empirically how hard would it be to break the exchange, *on average, on any random inputs drawn from the underlying group*. The groups under study were small in order (very far from the typical cryptographic groups used in practice), but we give compelling evidence that the security of the exchange tends to be dependent on the structure of the underlying groups. That structure can be recovered and rediscovered over and over as the group size increases.

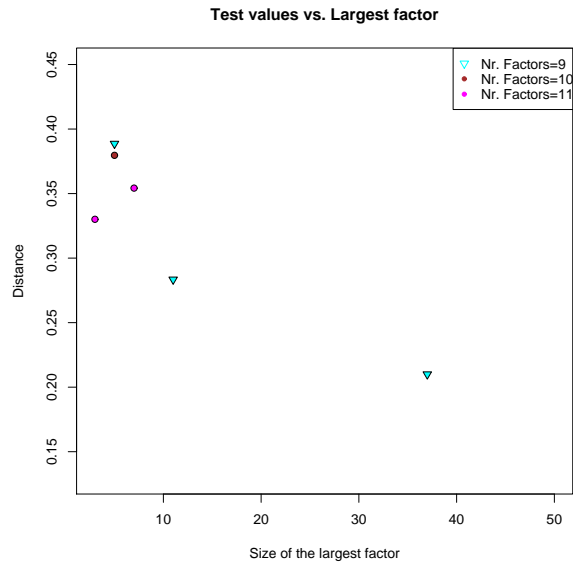
We have studied the relationship between the security of the Diffie Hellman public key exchange protocol and the structure of the underlying group. We first looked at groups where the protocol is provable not secure (working with cyclic groups of similar structure with multiplicative (\mathbb{Z}_p^*, \cdot)). We have found compelling evidence that breaking it (in the sense of actually finding the key) is dependent not only on the size of the largest factor in the decomposition of $p - 1$ but also on the number of terms in the decomposition. Furthermore, the relationship is not straightforward (as either one increases the security increases) since the interaction between these two determining factors is statistically significant. This means that it is entirely possible to have a group with large prime factor in the decomposition and a large number of terms in the decomposition of $p - 1$ and yet to be easier to break (on average for random inputs) than another groups where both these factors are smaller but they interact in a different way.

We show using statistical arguments that the prime subgroups of the (\mathbb{Z}_p^*, \cdot) are the most secure groups we have studied. We also show if one assumes that the structure of the group from which the subgroups are drawn remains the same, increasing the group's size translates into increasing the security of the Diffie-Hellman exchange as well.



(a) Test values for 3, 4, and 5 factors in the decomposition of $p - 1$

(b) Test values for 6, 7, and 8 factors in the decomposition of $p - 1$



(c) Test values for 9, 10, and 11 factors in the decomposition of $p - 1$

Fig. 4. If the two determining elements (number of factors and the size of the largest factor in the decomposition of $p - 1$) are independent we should see points of the same color close to parallel lines.




Fig. 5. We are comparing the prime subgroups with the corresponding safe groups. Values closer to zero represent safer groups for DH exchange.

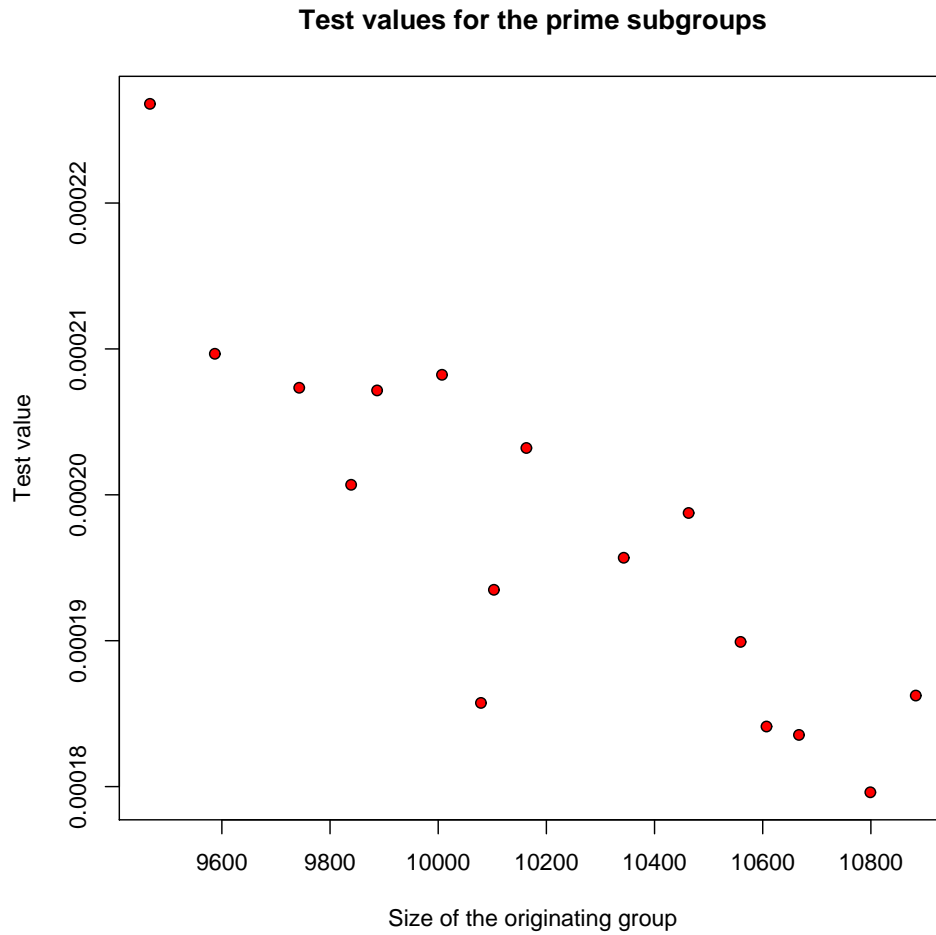


Fig. 6. Here we are showing that as the size of the group increases the values of the test do in fact decrease. Here we plot the corresponding values for the upper histogram in Figure 5. We note that the test values for this part were estimated using the same sample size (8 million) to insure that the test values are comparable and that variability of the test values is the same regardless of the size of the group.

Bibliography

- [1] D. Boneh. The Decision Diffie-Hellman problem. *Lecture Notes in Computer Science*, 1423:48–63, 1998.
- [2] Dan Boneh and Richard J. Lipton. Algorithms for black-box fields and their application to cryptography (extended abstract). In *CRYPTO '96: Proceedings of the 16th Annual International Cryptology Conference on Advances in Cryptology*, pages 283–297, London, UK, 1996. Springer-Verlag.
- [3] R. Canetti, J. Friedlander, and I. Shparlinski. On certain exponential sums and the distribution of Diffie-Hellman triples. *J. London Math. Soc.*, 59:799–812, 1999.
- [4] Ran Canetti, John Friedlander, Sergei Konyagin, Michael Larsen, Daniel Lieman, and Igor Shparlinski. On the statistical properties of Diffie-Hellman distributions. *Israel Journal of Mathematics*, 120(part A):23–46, 2000.
- [5] Ionuț Florescu, Alexey Myashnikov, and Ayan Mahalanobis. Statistical analysis of the diffie-hellman key exchange protocol in a finite group. Technical report, preprint available on ArXiv, 2007.
- [6] J. Friedlander and I. Shparlinski. On the distribution of Diffie-Hellman triples with sparse exponents. *SIAM Journal on Discrete Mathematics*, 14:162–169, 2001.
- [7] Rosario Gennaro, Hugo Krawczyk, and Tal Rabin. Secure hashed Diffie-Hellman over non-DDH groups. In *Advances in Cryptology - EUROCRYPT 2004*, Lecture Notes in Computer Science, pages 361–381. Springer Berlin / Heidelberg, 2004.
- [8] A. Joux and K. Nguyen. Separating Decision Diffie-Hellman from Computational Diffie-Hellman in cryptographic groups. *Journal of Cryptology*, 16:239–247, 2003.
- [9] S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, (22):79–86, 1951.
- [10] Ueli M. Maurer and Stefan Wolf. The relationship between breaking the Diffie-Hellman protocol and computing Discrete Logarithms. *SIAM J. Comput.*, 28(5):1689–1721, 1999.
- [11] M. I. Gonzalez Vasco, M. Näslund, and I. Shparlinski. New results on the hardness of Diffie-Hellman bits. In *Proc. Intern. Workshop on Public Key Cryptography*, volume 2947 of *Lect. Notes in Comp. Sci.*, pages 159–172, Singapore, 2004. Springer-Verlag.