

Unicity distance of the Zodiac-340 cipher

Joachim von zur Gathen
Universität Bonn, Germany
gathen@bit.uni-bonn.de

Abstract

In December 2020, David Oranchak, Jarl Van Eycke, and Sam Blake solved a 51-year old mystery: the *Zodiac cipher* of 340 symbols. Blake (2021) explains their solution. The correctness of their solution has not been seriously doubted, and here we give a further argument in its favor: the unicity distance of the cipher's system is between 80 and 152.

1 Introduction

In 1968 and 1969, a serial murderer killed five people in the San Francisco Bay area. He called himself Zodiac and bragged about his feats in several letters to local Police Departments and newspapers. Some of them were encrypted, one with 408 and another one with 340 symbols. They are now called Zodiac-408 and Zodiac-340, respectively. Some other coded texts are too short to allow deciphering. More murders and other messages have been connected to Zodiac, but these are not confirmed. In spite of the many clues he provided, the criminal has never been identified.

Zodiac-408 uses a homophonic substitution and was solved within a week by teacher Donald Harden and his wife Bettye. But Zodiac-340, mailed on a postcard on 8 November 1969, remained a major challenge to codebreakers. Klaus Schmech's blogpost (2020) features it as the second-most important unsolved cryptogram, after the Voynich manuscript.

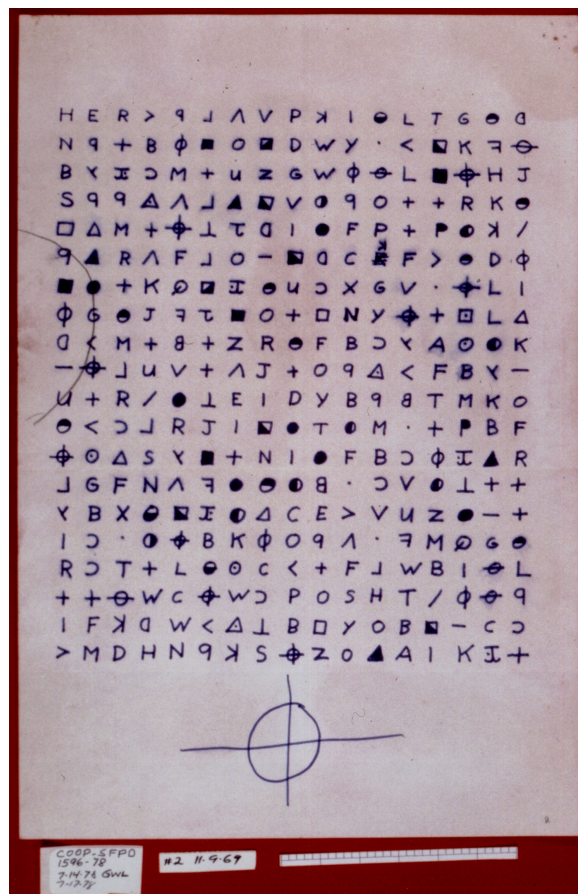


Figure 1: Zodiac-340 ciphertext.¹

Many people from all stations of life were attracted by this challenge, which is stated in an attractively concise form. Indeed, several solutions have been proposed, but none of them convinced the majority of experts. Do we need the sophisticated math and massive computing power of modern cryptology?

Yes, we do. American software engineer David Oranchak started in March 2013 the

¹<https://commons.wikimedia.org/w/index.php?curid=75983993>, last accessed 30 September 2021.

website <http://zodiackiller.net> which organized the efforts on Zodiac-340 in a systematic way, both human ingenuity and computing power, observations by interested people and software projects. This crowd-thinking and crowd-computing project bore fruit on 11 December 2020 when Oranchak, together with Australian mathematician Sam Blake and Belgian programmer Jarl Van Eycke, announced a break of the cryptogram. Their three talks Oranchak (2021a), Blake (2021), and Van Eycke (2021) formed the highly applauded key note address at the HistoCrypt conference in 2021. The present paper now calls Blake, Oranchak, and Van Eycke together the *Zodiac breakers*.

The correctness of their solution has not been seriously challenged and was publicly confirmed by the FBI. The present work shows that also Shannon’s theory of unicity distance in deciphering supports the solution.

The ciphertext is shown in Figure 1 and the plaintext reads as follows:

I HOPE YOU ARE HAVING LOTS
 OF FAN IN TRYING TO CATCH
 ME THAT WASNT ME ON THE
 TV SHOW WHICH BRINGO UP A
 POINT ABOUT ME I AM NOT
 AFRAID OF THE GAS CHAMBER
 BECAASE IT WILL SEND ME TO
 PARADLCE ALL THE SOOHER BE-
 CAUSE E NOW HAVE ENOUGH
 SLAVES TO WORV FOR ME WHERE
 EVERYONE ELSE HAS NOTHING
 WHEN THEY REACH PARADICE
 SO THEY ARE AFRAID OF DEATH
 I AM NOT AFRAID BECAUSE I
 VNOW THAT MY NEW LIFE IS
 LIFE WILL BE AN EASY ONE IN
 PARADICE DEATH

Blanks have been inserted appropriately, but no other changes were made. In particular, obvious original typos have not been corrected. California used gas chambers at that time to execute the capital punishment.

2 Unicity distance

Among the many solutions of Zodiac-340 that were proposed, which one is a “better” one, or “the correct” one? People will hold different opinions, in particular, the solvers about their own solution.

But there is a scientific answer to this question, based on Shannon’s theory of unicity distance. It requires the description of a system of encryption using a secret key, of which a specific instance is used. Then it yields a certain value, the *unicity distance*, so that any decipherment of a text which is longer than this value is highly likely to be unique and, within this theory, is accepted as correct.

The goal of this text is to provide such a system for Zodiac-340 and to analyze it. The conclusion is that the solution given above is correct. To the author’s knowledge, no such system has been put forth for any other proposed solution.

The American mathematician, electrical engineer and cryptographer Claude Elwood Shannon (1916-2001) laid the information-theoretic foundations of communication and cryptography in two papers from 1948 and 1949. He defined notions of information entropy and information content on probability spaces. Of interest to us is his notion of the *unicity distance* d :

$$d = \frac{I(\text{key})}{\log_2(\text{len}) - H(\text{lang})}. \quad (1)$$

This applies to the deciphering of a text of len many bits, encoded in a block cipher system with keys of information content $I(\text{key})$ bits, where the cleartext comes from a language with entropy $H(\text{lang})$ and \log_2 is the (binary) logarithm in base 2. Shannon’s famous theorem asserts that when the ciphertext has more than d symbols, then the decipherment is expected to be unique. We now view the Zodiac system as a method for encrypting 340-letter messages of the type that the killer sent.

As a side remark, Shannon’s information-theoretic approach remains

valid today and we employ it here. However, it is now largely replaced by a complexity-theoretic approach, since the fundamental paper of Diffie & Hellman (1976) that founded modern cryptology. In particular, it allows the exchange of secret keys over public channels, which is impossible information-theoretically.

The Zodiac solvers have discovered a method by which the cryptogram might have been derived. Formalizing their findings provides a system to which we may apply Shannon’s approach *mutatis mutandis*. To this end, we study the various contributions to the key space in Sections 3 through 7, then the language entropy in Sections 8 and 9, and derive an upper bound on the unicity distance in Section 8. Finally, we add some remarks on alternatives to the present approach.

Reichmann (2021) also argues for the correctness of the solution, mentioning “unicity distance”.

3 Homophonic substitutions

The entropy of random choices plays a central role in Shannon’s theory. Its simplest version refers to a finite probability space A whose elements a are equipped with a nonnegative probability p_a of occurring. A condition is that $\sum_{a \in A} p_a = 1$. Then the entropy is

$$H = - \sum_{a \in A} p_a \log_2(p_a). \quad (2)$$

The minus sign is required because $\log_2(p_a)$ is never positive.

This measure is appropriate in some cases, for example, for a uniformly random choice of keys among S possibilities. Then $p_a = 1/S$ for all keys a , each summand in (2) is equal to $1/S \cdot \log_2(1/S) = -\log_2(S)/S$. There are S summands, and taking into account the minus sign, we obtain $H = \log_2(S)$.

Zodiac-408 uses a homophonic substitution, and so does Zodiac-340. It is used for

encoding 26 letters of English in 63 symbols of Zodiac’s invention. Its choice contributes a large part to the keys’ security.

In general, we have two finite sets (alphabets) X of m plaintext letters (or words) and Y of n ciphertext symbols and associate to each plaintext letter some ciphertext symbols, also maybe none. Mathematically, it is not in general a function from X to Y , but a function $f: Y \rightarrow X$. This f corresponds to the decryption step, whose result is assumed to be unique.

The number of all such functions f is m^n . Thus the key space for these substitutions consists of exactly 26^{63} elements, and the information content of a key chosen uniformly at random is

$$I(\text{subs}) = \log_2(26^{63}) \approx 296.13. \quad (3)$$

In fact, this is an overestimate for cryptographic homophonic substitutions where the number of homophones will, at least roughly, follow the distribution of letters. Dhavare, Low & Stamp (2013) use a different description and find, in their equation (1), a value which is over 2^{24} times larger than the one above.

4 Sectioned plaintext

The cryptogram consists of 20 rows, each with 17 symbols. In the course of their work, the Zodiac breakers suspected (correctly) that the plaintext might have been divided into several sections. They tried 1 to 3 horizontal sections, each consisting of contiguous horizontal rows among the 20 rows in the ciphertext, and similarly for vertical sections of the 17 columns.

If the horizontal sections contain r_1, r_2, r_3 contiguous rows, with nonnegative values r_i and $r_1 + r_2 + r_3 = 20$, then these numbers form an *composition of 20 into at most 3 parts*. The number of compositions of an integer m into exactly i parts is $\binom{m-1}{i-1}$, and so the number of possibilities for horizontal

and vertical sections is

$$\left(\sum_{1 \leq k \leq 3} \binom{19}{k-1}\right) \cdot \left(\sum_{1 \leq k \leq 3} \binom{16}{k-1}\right) = 191 \cdot 137 = 26\,167. \quad (4)$$

Thus the entropy contribution of sectioning is

$$I(\text{sect}) = \log_2(26\,167) \approx 14.68. \quad (5)$$

5 Transpositions

Transpositions are a further classical tool in cryptography. In general, the plaintext is presented as a string $x_0, x_1, x_2, \dots, x_{m-1}$ of m symbols and a transposition length t is chosen. Starting with $y_0 = x_0$, every t th letter of x occurs in the transposed text y :

$$(y_0, y_1, y_2, \dots, y_{m-1}) = (x_0, x_t, x_{2t}, \dots, x_{(m-1)t}),$$

where the indices of the x 's are taken modulo m . This works if m and t are coprime, and then $y_j = x_i$ with $j \equiv it \pmod{m}$. In the decryption step, the y_j are given and the same relation between i and j now reads, equivalently, as $i \equiv jt^{-1} \pmod{m}$, where t^{-1} is the modular inverse of t modulo m . Implicitly, this involves a wrap-around: after the last entry comes the first one. The string is not a straight segment, but considered as a ring where the two ends of the segment are glued together.

This is a purely syntactical operation on indices, which does not depend on the values (or meanings) of the symbols. In contrast to homophonic substitutions, letter frequencies are unchanged, but digram frequencies may differ substantially.

As an example with $m = 9 \cdot 17 = 153$ and $t = 19$, the index sequences start with

$$\begin{array}{cccccccc} 0 & 1 & 2 & 3 & 4 & 5 & 6 & \dots \\ 0 & 19 & 38 & 57 & 76 & 95 & 114 & \dots \\ \dots & 7 & 8 & 9 & & & & \dots \text{ for } x, \\ \dots & 133 & 152 & 171 \equiv 18 & & & & \dots \text{ for } y, \end{array}$$

so that the transposition of $x_0, x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9, \dots$ is

$$\begin{aligned} & y_0, y_1, y_2, y_3, y_4, y_5, y_6, y_7, y_8, y_9, \dots \\ = & x_0, x_{19}, x_{38}, x_{57}, x_{76}, x_{95}, x_{114}, x_{133}, \\ & x_{152}, x_{18}, \dots \end{aligned}$$

For the last entry, we note that $-8 \cdot 19 = -152 = -153 + 1 \equiv 1 \pmod{153}$ and therefore $t^{-1} = 19^{-1} \equiv -8 \equiv 145 \pmod{153}$. Indeed, $18 \equiv 9 \cdot 19 \pmod{153}$, and conversely $9 \equiv 18 \cdot 145 \pmod{153}$.

In the above, the plaintext is given as a string in a one-dimensional format. The Zodiac cipher uses a two-dimensional variant of this, of which no other example seems to be known. Figure 2, produced by Sam Blake, shows how this is applied to the top 9 (of 20) rows. The numbering of rows is $0, 1, \dots, 8$ and that of columns is $0, 1, \dots, 16$ in the Zodiac text. The second entry above says that $y_1 = x_{19}$ and can be seen in the yellow box in the second row and third column. Similarly, $y_8 = x_{152}$ is visualized by the green-yellow 8 in the lower right corner. That field bears the largest index for the x 's, namely 152.

Since the transposition length 19 is larger by 2 than the rectangle width of 17, adding 19 corresponds to moving 2 steps to the right and 1 step down. This *2-1-move* or knight's move is clearly visible in Figure 2.




Figure 2: Numbered top 9 rows of Zodiac-340 transposed.

So far, so good. What comes after the last entry? Of course, the first one. So the rectangle turns into a donut, where the left and right sides as well as top and bottom are glued together. Where does a knight's move take us from the lower right 8? Two to the right, into the position numbered with 9, and then 1 down, to the box with 145. But the Zodiac scheme actually takes us to the 9.

This happens systematically. At the right-hand edge, a knight's move takes us

2 to the right, in the next row, and then 1 down. But the down step is not taken, instead only a 2-0-move. This avoids leaving a row unused at that point, and we call this procedure *no unused rows*.

With this transposition, the first nine rows of the cryptogram decipher as given on page 2. The whole text is split into three horizontal rectangles, all of 17 columns and of 9, 9, and 2 rows, respectively. The middle and bottom rectangles are shown in Figure 3, also by Sam Blake.

153	162	171	180	189	198	207	216	225	234	243	300	301	302	303	304	305
284	292	154	163	172	181	190	199	208	217	226	235	244	252	260	268	276
269	277	285	293	155	164	173	182	191	200	209	218	227	236	245	253	261
254	262	270	278	286	294	156	165	174	183	192	201	210	219	228	237	246
238	247	255	263	271	279	287	295	157	166	175	184	193	202	211	220	229
221	230	239	256	264	272	280	288	296	158	167	176	185	194	203	212	248
204	213	222	231	240	249	257	265	273	281	289	297	159	168	177	186	195
187	196	205	214	223	232	241	250	258	266	274	282	290	298	160	169	178
170	179	188	197	206	215	224	233	242	251	259	267	275	283	291	299	161
309	308	307	306	310	311	312	313	315	314	317	316	318	319	320	321	324
323	322	326	325	334	333	332	331	330	329	328	327	335	336	337	338	339

Figure 3: Middle and bottom sections of Zodiac-340 transposed.

The middle section looks pretty much like Figure 2, but with two modifications. The last six entries in the first row correspond to the cleartext LIFEIS and do not participate in the 2-1-transposition. In the sixth row, the last entry labelled 248 has been moved from its proper position in the fourth column (which is now labelled 256) to the last column.

The bottom section of two rows does not involve any transposition. Of its nine words, three are spelled correctly (increasing numbers) and six are written backwards (decreasing numbers). Reversing words does not create a big problem for human or machine decipherers and thus does not contribute much to security. We ignore it in the following except that we grant one bit for “use word reversals” or not.

Irregularities in a stepping function can make a cipher substantially more secure. As a principle, this was employed in the

German cipher machines *Lorenz Schlüsselzusatz SZ-40* during the Second World War, and in the Swiss version *NEMA* of the *Enigma*, built just after that war. In fact, the irregularities in the Zodiac-340 transposition posed a serious difficulty for the breakers.

We now work with an arbitrary transposition length from 0 (no transposition) to 51 and the two-bit choice to use *no unused rows* and *word reversals* or not. This gives $52 \cdot 2 \cdot 2 = 208$ as the total number of possibilities and the contribution to $I(\text{key})$:

$$I(\text{trans}) = \log_2(208) \approx 7.70. \quad (6)$$

6 Irregular substitutions

Some aspects of the Zodiac-340 cryptogram are not captured by the above considerations on homophonic substitutions, sectioning, and transposition. These are:

- Misspellings.**
- Five words are misspelled: FAN, BRINGO, BECAASE, SOOHER, E for FUN, BRINGS, BECAUSE, SOONER, I.
 - The Zodiac-340 substitution has no ciphertext symbol for K (in contrast to Zodiac-408), and the two occurrences of the letter K are written as V: WORV, VNOW for WORK, KNOW.
 - The incorrect PARADICE appears already in Zodiac-408, elsewhere in the Zodiac corpus, and three times in Zodiac-340, once even further contorted with a typo as PARADLCE.

Nulls. The text LIFEIS in the penultimate row of the cleartext consists of nulls (dummies), maybe serving to make the text fit exactly into its rectangular array.

Skip. In row 15 (sixth row of the middle section) a letter is moved from its proper position in the fourth column

to the last column in the same row. This is the box labelled 248 in Figure 3.

All these can be described by allowing the following type of *replacement* in the encryption. We augment fictitiously the 26-letter English alphabet by one more, the empty space \star . This is not a blank, but an invisible character. Thus BRINGO and BRIN \star GO read exactly the same way. Replacing a letter by \star means removing that letter. Then all of the modifications listed above can be described as picking a position in the plaintext and replacing the character at that position by one of those 27 symbols. A skip corresponds to two replacements, although these are special in that the deleted and the inserted letter are the same. Another special case would be to not change anything, emulated as replacing a letter by itself.

operation	number	replacements
misspelling	5	5
V for K	2	2
PARAD(I/L)CE	3	4
dummy	6	6
skips	2	4
total		21

Thus we have a total of 21 replacements in a 340-letter text, which comes to about 6.18%. If we allow generously 25 replacements, then there are $R = \binom{340}{25} \cdot 27^{25}$ possibilities, with a contribution to $I(\text{key})$ of

$$I(\text{replace}) = \log_2(R) \approx 244.12. \quad (7)$$

7 Key entropy and text length

We are now ready to determine the value of $I(\text{key})$ in (1). A key consists of several parts, each of which is chosen uniformly at random and independently. The rounded values are:

Homophonic substitution. $I(\text{subs}) = \log_2(26^{63}) \approx 296.13$, by (3).

Sectioning. $I(\text{sect}) = \log_2(26 \cdot 167) \approx 14.68$. as given in (4).

Transpositions and word reversals.

$$I(\text{trans}) = \log_2(208) \approx 7.7, \text{ by (6).}$$

Replacements. $I(\text{replace}) = \log_2\left(\binom{340}{25} \cdot 27^{25}\right) \approx 244.12$ by (7).

Total $I(\text{key}) \approx 562.62$ by adding up the four contributions above.

In Shannon's theory, ciphertext symbols are supposed to be uniformly distributed, and we now assume this to be the case for the Zodiac cipher. This is consistent with the fact that, before its solution, the possibility that it might be gibberish has been seriously considered by many; see Oranchak (2018). Then the information content of a ciphertext of k symbols is $k \log_2 63 \approx 5.98 k$ bits. For Zodiac-340, this comes to $340 \cdot \log_2(63) \approx 2032.28$ bits, and in (1), we have

$$\log_2(\text{len}) = \log_2(2032.28) \approx 10.99 \approx 11. \quad (8)$$

In (1), the language entropy $H(\text{lang})$ does not refer to the ciphertext, but to the plaintext of the cryptogram, and is not directly related to the deciphering effort. We first have to determine the length of the plaintext. One might be tempted to assume it as 340 letters, but that is not correct.

Any claimed solution that somehow substitutes and rearranges the cipher symbols will be a single word of 340 letters and certainly not an English text. It (almost) becomes one if we insert 90 blanks appropriately, as done on page 2. Thus the plaintext consists of 430 characters in a 27-letter alphabet, including the blank. In general, the average English word length is estimated at 4.5 nonblank letters; see Shannon (1951), Section 2. Thus an English text of ℓ characters can be expected to reduce to $k = \ell(1 - 1/4.5)$ letters when the blanks are removed. In other words, a reduced text of k letters corresponds to a regular text of $\ell = 9k/7$ letters. And indeed, the fraction $9/7 \approx 1.286$ matches quite well our value of $430/340 \approx 1.265$. Thus we will take

$k \cdot 430/340$ letters as the length of a plaintext encrypted by k symbols. We could skip this step if the relevant papers cited here also excluded the blanks—but they do not.

8 Language entropy

The only ingredient to (1) still missing is the language entropy $H(\text{lang})$. For a complicated probability space, say, texts in a natural language such as English, a naïve application of (2) fails to be the appropriate measure. It only takes into account the frequency distribution on individual letters and is called the *monogram* (or *single-letter* or *1-gram*) entropy. It evaluates to about 4.1, and one often sees such incorrect values in some parts of the literature. Also other issues around this entropy are often not properly taken into account. In particular, sometimes the most frequent character in English text is ignored: the blank \square .

The monogram entropy does not reflect the rich structure of English, where individual words and phrases also occur repeatedly. A basic reason is that the corpus of all English texts is not finite, and even fairly large but still finite compilations do not yield a reliable result. Longer *polygrams* (often called *n-grams* for some specific value of n) also have to be considered. For any value of n and a given text (of finite length), the entropy E_n of n -grams is calculated according to (2), and the *conditional entropy* of n -grams over $(n - 1)$ -grams is $F_n = E_n - E_{n-1}$. This F_n refers to the prediction of the next letter, when the previous $n - 1$ ones are known.

According to Shannon [10], Section 7, the sequence of F_n for growing n approximates the entropy, here of English, better and better. Unfortunately, these values are hard to compute. Shannon (1951), Section 6, calculates experimentally bounds for the F_n , for example, $1.3 \leq F_6 \leq 2.2$. Goldreich et al. (1999) show that under standard complexity-theoretic assumption, ar-

bitrarily good approximations are infeasible to compute. Experiments with a corpus of two billion characters in von zur Gathen & Loebenberg (2017), Figure 3, illustrate the practical issues: for monograms ($n = 1$) the value F_1 is 3 to 4 times too large, the F_n remain too large for n up to 4, they lie in Shannon’s interval for $5 \leq n \leq 11$, and are too low for larger n . The computation gets distorted by “noise”, since those longer n -grams do not have enough “room” to display their true frequencies.

Now we need to determine the plaintext entropy of the cryptogram’s plaintext language. One can consider (at least) three “languages” to give rise to the 430-letter plaintext:

- standard English,
- the language of the Zodiac-340 cryptogram, as given on page 2,
- the language of the Zodiac corpus.

For standard English, we may assume an entropy around 1.5, but see the provisos mentioned above. The Zodiac-340 cryptogram has an entropy around 1.8. This is calculated as for the Zodiac corpus in the next section and we forego the details. Dhavare et al. (2013) apply a hill climbing method to Zodiac-340, but do not obtain a satisfactory result. In their subsection 6.3, they point out “that our algorithm is very sensitive to the language model ... Therefore, it might be worthwhile to consider specialized “language” models when trying to break the Zodiac-340.”

In a different context, we now do exactly that and concentrate on the *Zodiac corpus*, consisting of 20 messages from the Zodiac killer, which date from 31 July 1969 to 8 July 1974; see Oranchak (2021b). Most of them were sent in plaintext to Californian newspapers and police departments, to a lawyer, and one scribbled on a victim’s car door. Also included are the plaintexts of the Zodiac-408 and Zodiac-340 cryptograms with blanks appropriately

inserted. Some Zodiac cryptograms that are too short to be deciphered and must be left out. There are also spurious messages whose claim to be from Zodiac is disputed.

Plaintexts of the Zodiac cryptograms do not contain numerals or punctuation marks and for this study, they were removed. The Zodiac corpus then contains 14859 letters and blanks. Its entropy may be estimated to be around 1.8; details are given in the next section.

However, this is much ado about nothing. Whichever approach from the three listed above we take, the entropy comes out to be between 1.3 and 2.3, and the unicity distance is only slightly sensitive to its exact value.

Shannon’s fundamental idea is that if the information content of the ciphertext is larger than that of keys and plaintext combined, then one can expect a unique deciphering solution. For a message of k symbols in the Zodiac system, the plaintext length is $k \cdot 430/340$ letters under the language distribution and Shannon’s condition is that $k \log_2(63) \geq k \cdot 430/340 H(\text{lang}) + I(\text{key})$. Rearranging and using $I(\text{key})$ from Section 7 and $H(\text{lang}) = 1.8$, this amounts to

$$k \geq \frac{I(\text{key})}{\log_2 63 - 430/340 H(\text{lang})} \approx \frac{562.62}{5.98 - 1.8 \cdot 430/340} \approx 152.03.$$

In order to be on the safe side, we have consistently overestimated various quantities. If we just consider the homophonic substitution part of the Zodiac-340 cipher and take the value 296.13 from (3) instead of the above $I(\text{key})$, then we find the lower bound 80.019 on the unicity distance.

The unicity distance of the Zodiac-340 cipher is between 80 and 152.

The actual length of 340 of the cryptogram is much larger than the upper bound.

9 Zodiac language entropy

Frequency calculations are an essential tool in cryptanalysis. In fact, the observation that a guessed transposition of 19 increases substantially the number of repeated digrams in the cryptogram was a vital step in the Zodiac break. However, the following calculations are not related to cryptanalysis, rather they concern frequencies in the 14825-character Zodiac plaintext corpus. It uses the 26 letters of the English alphabet and the blank \square . For n up to 5, we list the three most frequent n -grams, their number of occurrences and their rounded frequency in percent:

\square	E	T
2978, 20.08	1405, 9.47	1097, 7.39
$E\square$	$\square T$	TH
615, 4.14	489, 3.29	410, 2.76
$\square TH$	THE	$HE\square$
346, 2.33	274, 1.84	188, 1.26
$\square THE$	$THE\square$	ING \square
256, 1.72	182, 1.22	105, 0.70
$\square THE\square$	$\square YOU\square$	N \square THE
182, 1.22	49, 0.32	38, 0.25

We find the following entropies $\text{Ent}(n)$ and conditional entropies $\text{cEnt}(n) = \text{Ent}(n) - \text{Ent}(n-1)$, where we use $\text{Ent}(0) = 0$:

$n :$	1	2	3	4	5
$\text{Ent}(n)$	4.09	7.14	9.63	10.99	11.75
$\text{cEnt}(n)$	4.09	3.24	2.30	1.36	0.76

The noise discussed in Section 8 also distorts the pentagram conditional entropy here, and may affect the tetragram conditional entropy. We now take the mean of the tri- and tetragram conditional entropies as value, that is

$$H(\text{lang}) = 1.8. \tag{9}$$

Any such choice has an element of arbitrariness, as mentioned above. Below we illustrate the (limited) effect of this choice by also calculating with 1.36 and 2.30 as values of $H(\text{lang})$.

These values are at the upper end of or beyond the bounds that Shannon states.

Spelling rules make reading easier by increasing redundancy and thus reducing entropy. In fact, correcting 124 spelling mistakes in the Zodiac corpus changes the polygram entropies slightly, most notably $cEnt(3)$ from 2.30 to 2.17. On the other hand, $cEnt(4)$ increases slightly. This may indicate noise already for these tetragrams.

10 Alternatives

The estimates in Section 7 are taken rather generously and some may overshoot the real values substantially. That is acceptable, since it makes the final result on the unicity distance more reliable. We do not have the goal of lowering the estimate of the unicity distance to a more realistic value.

But if one wanted to, one might start with $I(\text{subs})$ in (3), an upper bound on the information content in Zodiac's homophonic substitution. Since that is supposed to level out frequencies, the rarely used letters in the English alphabet will have few (0 or 1) homophones and the frequent letters a higher number. These constitute only a small fraction of what we allow as key space for substitutions.

In a system for communicating secretly, a legitimate recipient in possession of the secret key can restore the plaintext correctly. This is not possible in the presence of spelling mistakes. So in such a system, one would ignore the option of making such errors, reducing the value of $I(\text{replace})$.

The language entropy is a fickle thing. For two values mentioned in Section 9, namely 1.36 and 2.30, we obtain unicity distances of 132.2 and 183.4, respectively. This shows robustness of the main claim under modified assumptions on the entropy.

11 Conclusion

The unicity distance for a ciphertext encrypted with a method as the Zodiac-340 cryptogram is between 80 and 152, under the assumptions stated above. The actual length 340 is much larger than this value.

Our findings show that the solution is correct beyond doubt.

The method includes four steps:

- A randomly chosen homophonic substitution of 26 letters in 63 symbols,
- a split into up to 3 horizontal or vertical sections,
- a transposition by up to 51 places, in the one-dimensional or the 2-1-dimensional sense,
- a certain number of arbitrary changes in individual letters, such as spelling mistakes.

Is there a different solution of the cryptogram? That is, can one come up with a well-specified system under which it could have been encrypted and whose unicity distance is below 340 (or below 152)?

Acknowledgements

This text would not exist without the success of the Zodiac solvers. In addition, they have contributed substantially with helpful discussions, hints, and suggestions.

Many thanks also go to Daniel Panario for his help with the composition(s) of this paper.

References

- Sam Blake. 2021. <https://www.youtube.com/watch?v=iuNyQ44JYxM>.
- Amrapali Dhavare, Richard M. Low, and Mark Stamp. 2013. Efficient cryptanalysis of homophonic substitution ciphers. *Cryptologia*, 37:250–281.
- Whitfield Diffie and Martin E. Hellman. 1976. New directions in cryptography. *IEEE Transaction on Information Theory*, IT-22(6):644–654, November. DOI 10.1109/TIT.1976.1055638.
- Joachim von zur Gathen and Daniel Loebenger. 2017. Why one cannot estimate the entropy of English by sampling. *Journal of Quantitative Linguistics*. 30 pages. <https://doi.org/10.1080/09296174.2017.1341724>.

- Oded Goldreich, Amit Sahai, and Salil Vadhan. 1999. Can statistical zero knowledge be made non-interactive? or On the relationship of SZK and $NISZK$. In *Springer Lecture Notes in Computer Science 1666*, pages 467–484.
- David Oranchak. 2018. Are the ciphers gibberish? <http://www.zodiackillerciphers.com/?p=774>.
- David Oranchak. 2021a. <https://www.youtube.com/watch?v=44rkCyU6ssE>.
- David Oranchak. 2021b. Zodiac killer letters. <https://github.com/doranchak/zodiac-killer-ciphers/tree/master/docs/letters>.
- F. Reichmann. 2021. Why the transposition in the 340 solution is inevitably correct. <https://zodiackiller.net/community/zodiac-cipher-mailings-discussion/why-the-transposition-in-the-340-solution-is-inevitably-correct/>. Last accessed 31 March 2022.
- Klaus Schmech. 2020. <https://scienceblogs.de/klausis-krypto-kolumne/the-top-50-unsolved-encrypted-messages/>.
- C. E. Shannon. 1948. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423 and 623–656. Reprinted in CLAUDE E. SHANNON and WARREN WEAVER, *The Mathematical Theory Of Communication*, University of Illinois Press, Urbana IL, 1949.
- C. E. Shannon. 1949. Communication theory of secrecy systems. *Bell System Technical Journal*, 28:656–715.
- C. E. Shannon. 1951. Prediction and entropy of printed English. *Bell System Technical Journal*, 30:50–64, January.
- Jarl Van Eycke. 2021. <https://docs.google.com/presentation/d/19PT51INr31jh9KL0oxlcA29pjJ2i9QfsipueS8PM0/edit?usp=sharing>.