

1999

Unclassified

APPROVED FOR
RELEASE DATE:
07-Jan-2009

Next

Previous

Contents

The Puzzle at CIA Headquarters

Cracking the Courtyard Crypto

(b)(1)
(b)(3)

Editor's Note: In the late 1980s, under a GSA program, the CIA commissioned Washington artist James Sanborn to create a series of sculptures for CIA's new Headquarters building. Working together with [redacted] who was soon to retire from the Office of Communications, Sanborn created a sculpture named "Kryptos" (Greek for hidden) that was dedicated in 1990 and now resides in the northwest corner of the courtyard. The curving verdigris scroll contains an 865-character coded message that seems to flow out from a petrified tree and is located near a water-filled basin bordered by various types of stones. In the following article, [redacted] describes how he has deciphered most of the secret message contained within the sculpture.

I looked down at the large number of yellow pages littering the top of my desk, obscuring its dark oak finish. Each page was covered in hundreds of alphabetic letters that, although they appeared to be arranged randomly, were really carefully laid out in columns of hidden order. For the past week or so, I had been feeling that I was very close to a solution--but was I? As a hobby, I had solved many puzzles since I was a child, and I had discovered that the most difficult challenge was continually dealing with the nagging fear that I may be on the wrong track. The voice in my head would whisper, "Maybe you're mistaken--maybe this first part of the Kryptos code is really not a polyalphabetic Vigenere Tableau after all--maybe it's a different type of code entirely. Or maybe it is a Vigenere code, but it's been double or triple encoded--or maybe it was encoded backwards, or maybe..."

This is the worst train of thought an amateur cryptanalyst can have--it goes nowhere, and it is tempting to give in to frustration at the seemingly infinite number of possible solutions and abandon the effort. So, as I had done dozens of times before, I pushed these thoughts away and pressed on. The thing to do, I told myself again and again, was simply to stick to what I knew to be the facts. Let the encoded text itself guide me, and keep speculation to an absolute minimum. I should follow my instincts occasionally, but then remember to always keep track of my assumptions and be prepared to revise my thinking if a hunch doesn't work out.

And suddenly it happened--I was hit by that sweetly ecstatic, rare experience that I have heard described as a "moment of clarity." All the doubts and speculations about the thousands of possible alternate paths simply melted away, and I clearly saw the one correct course laid out in front of me. Taking a fresh sheet of paper, I slowly and deliberately wrote out a new column of letters, followed by another, and then another. I continued this for several pages, then computed mathematically which rows were most likely to represent the correct plaintext letters, and searched for logical combinations between adjacent letters. I tried to contain my excitement as I witnessed the miracle of letters slowly forming together into words, one after the other. Within the next few hours, I had finished. After more than seven years encompassing some 400 hours of laboring over piles and piles of paper covered with gibberish, I was at last looking down at a paragraph of clear English text. I had

broken out the first part of the Kryptos code.

A New Challenge

I don't remember the first time I saw, or even heard about, the Kryptos sculpture in the courtyard. But each day as I walked past the huge, curved copper plates, I would look at their mysterious perforated inscriptions a little longer each time until gradually I was lured into an attempt to decrypt them. Like the character of Ishmael from Melville's *Moby Dick*, I, too, embark on journeys "... whenever I find myself growing grim about the mouth" or find my soul becoming "... damp, drizzly November." Rather than putting to sea as Ishmael would have done, however, my travels are journeys of the mind. By undertaking new challenges, I explore the mysteries at hand.

The author and the courtyard crypto. CIA photo.

I quickly became fascinated with the thought of that encoded message standing mutely in the Headquarters' courtyard year after year, wordlessly taunting everyone to try to read its hidden message; indeed, at its dedication, the puzzle was offered to Agency employees as a challenge. I have always been interested in all different types of puzzles: I've spent weeks uncoupling bits of twisted metal that some sadist has carefully coupled; I've poured over endless permutations of logic problems that cheerfully describe groups of people with terribly contrived lives who, it seems, are constrained to sit at the dinner table only in unique arrangements; and I've patiently studied colored cubes, labyrinths, cryptograms, and pictures that have been jigsawed into thousands of identical pieces. But I had never before undertaken a puzzle anything like that represented in the Kryptos sculpture.

Because I'm not a professional cryptanalyst, I first needed to educate myself in the different techniques used in the field of cryptography. After working on the Kryptos codes for awhile, I soon learned that the basic problem-solving techniques that I needed for decryption were not so different from those that I used in my work as a technical analyst. And in my reading I also learned that codebreaking is more than just a fascinating hobby; throughout history, lives literally have been lost and saved due to cryptanalysis. It has been estimated that, for example, the war in the Pacific in World War II was shortened by at least a year as a result of the efforts of military cryptanalysts.

From the time I initially began sharing with people the progress I had made deciphering the code covering the Kryptos sculpture, I have had the interesting experience of being inundated with questions about it. By far the most frequent question asked is, "How much did you get paid for solving that thing?" (Many people seemed to think that there was some kind of prize involved.) The next most common question is, "How long did it take you to solve it?" Notwithstanding my wife's dry riposte ("So long, that it almost wasn't worth it"), this question is more difficult. I was not keeping track of the time, and, in any event, I was not in any hurry to solve the puzzle. For me, when codebreaking ceases to be fun and begins to seem more like work, then it's time to quit and start doing something else. My best estimate of roughly 400 hours seems about right, had I been working continuously.

I had expected the most prevalent question to be, "What does the message say?"; this one, however, only made it into third place, followed closely by (usually humorously) "Don't you have anything better to do with your time?" Although I'm not sure just how to respond to this last one, I do like to talk about the progress I've made on the Kryptos problem. But I must confess that I am often hesitant to reveal the message without having an opportunity to provide a full explanation.

There are several reasons that I don't like to just blurt out the solution to the Kryptos puzzle. First, it is important to remember that the message is still incomplete; I have not yet broken out the last 97 characters of the 865-character inscription. Second, the message is enigmatic and open to interpretation. Finally, I believe that simply presenting the solution without showing the methodology behind it is cheating people a little. The creators of Kryptos no doubt intended to inspire people to try to solve it for themselves. Merely showing the "solution" that I have derived steals away a little of the excitement and appreciation of the problem, and can discourage people from trying their own interpretations.

When asked to write an article for *Studies in Intelligence*, I became intrigued by the idea of trying to explain what can be a complicated and convoluted subject in a manner that could be understood by anyone. After it was suggested to me that it might well be impossible to explain the methodology concisely and in a nontechnical manner, I couldn't resist taking up the challenge. The purpose of this article is not simply to reveal the content of the Kryptos message but rather to give the reader a sense of the exhilaration that is felt when solving complex problems after struggling with them for years.

Where To Begin?

When first confronted with the coded full text of the Kryptos sculpture, the task of deciphering it can seem formidable--even impossible. Even the most casual of observations, however, reveals that the writing on the sculpture is divided into two main pieces. One side (Figure 1) depicts a somewhat modified "Vigenere Tableau" (a series of alphabets used for the coding and decoding of text); the other side contains the actual coded message of 865 characters interspersed with four "question marks" (Figure 2).

One method for deciphering a code begins by counting up the number of times that each letter of the alphabet occurs throughout the message, a so-called frequency count. I saw no reason, however, why different sections of the Kryptos code could not have been enciphered with different coding schemes; if so, it would be necessary to perform frequency counts on each part separately. Looking through the code, I seemed to see different patterns throughout, but I wanted to quantify these impressions. Therefore, just as a starting point, I decided to count up the number of times each letter occurred for each row of the message depicted in Figure 2. By plotting these counts as a function of row for each letter of the alphabet (26 plots in all), I would be able to watch how often each letter was being used as the message evolved. In particular, if the count stayed approximately constant for a number of rows, this would indicate the same type of code was being used for that part of the message.

Figure 1. One side of Kryptos sculpture, depicting modified Vigenere Tableau.

Figure 2. Opposite side of Kryptos sculpture, containing the coded message.

"Encryption Schemes"

A typical example of one of the 26 plots, for encoded letter "J," is shown in Figure 3. (Of course, these "J"s are not really "J"s at all, but are ciphertext letters representing plaintext letters of the alphabet.) It can be seen that the letter "J" occurred twice in the first row of the Kryptos message, once in the second row, and so on. I noted that there was a marked change in letter frequency occurring between rows 14 and 15; that is, from row 1 to 14 there were "J"s appearing in almost every row, but from row 15 to row 25 there were no "J"s appearing at all. This indicated to me that rows 1-14 were encoded with a different type of code than rows 15 and beyond. In addition, it appeared that the last two or three rows were encoded with a different type of code than the preceding rows. The dramatic change from row 14 to 15 was seen in virtually all 26 frequency plots.

Figure 3. Frequency of "J" occurrences by row.

Because of these results, I divided the code into a top section (rows 1-14) and a bottom section (rows 15-28). I then further divided the top section into three parts using the four question marks as the dividing points; I designated these portions as Parts I, II, and III, respectively (Figure 4). I began my attempts at deciphering with the 125 characters of Part II mainly because the repetition of the "DQM" sequence occurring in the first 11 letters made this piece look more interesting to me than the other two parts.

Figure 4. Top section of Kryptos code, divided into three parts.

Frequency Counts

The frequency count is one of the cryptanalyst's most powerful tools. By computing how often each

letter appears throughout the message, one can determine important clues about the code employed. Figure 5 shows a typical frequency count of normal English text.

Figure 5. Normal frequency distribution for the English language.

Each bar represents the number of times that that letter occurs, on the average, for every 1,000 letters encountered. This pattern adheres so tenaciously to our language that it is virtually impossible to write or speak without it betraying its presence. And even if one were to artificially construct a few sentences or paragraphs that do not follow this distribution, it would turn out to be so contrived that new patterns would inevitably develop. Because it is an essential and unavoidable part of our language, from a cryptanalytic point of view, the pattern is also a code's greatest weakness.

In Figure 6, a frequency count for Part II of the Kryptos code is plotted. It is easily seen, from comparison with Figure 5, that the pattern is very different from the normal distribution for English.

Figure 6. Frequency count for Part II of Kryptos code.

In particular, the distribution appears much flatter. This is exactly what I would expect to see for a polyalphabetic substitution code, where the distributions for several different alphabets would tend to average out when combined. Because the modified Vigenere Tableau--a polyalphabetic substitution code--is inscribed on the Kryptos sculpture, it made sense to try this first as a working hypothesis for this part of the code.

Finding the Length of the Keyword

A common method for breaking a Vigenere code involves first determining the length of the keyword that was used to encipher the message. The code then can be broken up into its constituent single alphabets, called monoalphabets, and analysis performed on each cipher alphabet to determine to which Vigenere alphabet it corresponds. Once these Vigenere alphabets are determined, the code can be easily deciphered. Note that nowhere in this analysis is it necessary to actually resolve the content of the keyword.

To ascertain the length of the keyword used in Part II of the cipher, I needed to make use of a cryptanalytic tool known as the "index of coincidence," or I.C. This concept conveys the probability that, given a distribution of letters, any two chosen at random will be the same. The details are covered in the field of statistics, and they are not significant here; the formula and an example calculation is provided, (see Box B), if anyone wants to repeat the analysis. The important point is that for any set of letters, the average value of the I.C. will range from 0.038 up to 0.066. The closer this average value is to 0.038, the more likely it is that the set of letters is completely random. If the value is closer to 0.066, then it is more likely that the letters represent a monoalphabet.

Example Computation of I.C. for 7 Alphabets (Kryptos Code Part II)

I first divided the text of Part II into two single alphabets (corresponding to a two-letter keyword) and computed the average I.C., then divided the text into three alphabets and again computed the average I.C., then divided into four alphabets, and so on (see Box B for an example calculation using seven alphabets.) Finally, I plotted each average I.C. as a function of the number of alphabets (Figure 7).

It is easy to see from this figure that the correct key length for this piece of code should be eight because that number yields an average I.C. value of almost 0.063; thus, I knew that I needed to divide the code into eight separate alphabets (Figure 8).

Figure 7. I.C. vs. Number of Alphabets for Part II.

Figure 8. Part II of Kryptos code divided into eight alphabets.

Determining Which Vigenere Alphabets To Use

Now that I had the code divided into eight monoalphabets, I needed to determine which Vigenere alphabet corresponded to each of them. If the text of Part II had been much longer, the analysis would have been fairly straightforward; I could simply have performed a frequency count for each column and matched the resulting distributions to the English frequency count in Figure 5. The letter occurring most frequently would likely have corresponded to plaintext "E," and, because the Vigenere alphabets have a fixed order (cyclical permutation), the identities of the other letters could then have been found. From the example presented in Figure 9 (for the first column), however, it is obvious how impossible this method would be for the present problem:

Figure 9. Frequency count for first alphabet of Part II.

With only around 15 letters available per alphabet, there are just not enough to determine any kind of alphabetic structure such as that seen in Figure 5. So now what?

I must admit that I got stuck at this point for many months. No doubt a professional cryptanalyst would know several ways to proceed, but I was unsure. And then on February 21st, 1998, I had a breakthrough. How well I remember sitting at my desk that day, with the sunlight beaming through the window and the birds chirping unusually loudly outside. Suddenly, I felt a burst of Divine Inspiration--the insights I had gained about the Kryptos code from hundreds of hours of work came together and combined with all I had learned about cryptography. Everything seemed so clear that I didn't have the slightest doubt about how to proceed, and I was confident that my plan would work.

Once again I divided the message into eight separate alphabets. What had become plain to me was that if I wrote out one of the columns along with every possible permutation for each letter in that column, then exactly one of the resulting columns would contain the true plaintext letters. True, only every eighth letter of the message would be revealed (the first, ninth, seventeenth, and so forth), but I could similarly write out the permutations for every other column and then look for likely letter combinations among the plaintext (Figure 10).

One of these columns had to correspond to the correct plaintext letters for the first column of the message in Figure 8, but which one?

Looking across the columns (Figure 11), it can be seen that some appear less promising than others; those that contain a large number of uncommon letters (B, J, K, Q, X, Z), for instance (remember, these should now represent actual plaintext letters.) But this all seemed so subjective--and how could I tell which column was the best choice? Columns 1, 5, 7, 11, 15, and 16 each contain at least five uncommon letters--very unlikely to occur in such a short message of only 125 characters, so I could rule these out. But what about the others? Column 18, for example, contains four "R"s, which is a very common letter, and only one "J" (an uncommon letter). But is that column more likely to represent the plaintext letters than column 22, which has two "E"s and four "O"s, and no "K"s, "J"s, or "Z"s?

Figure 10. First column of Part II with every possible permutation.

Figure 11. Permutations of first column, with uncommon letters in red.

I needed to quantify these observations mathematically. One way to do this would be, for each column, to add up the normal English frequency count associated with each letter. Then the columns with the highest sums would be the best candidates to correspond with the true plaintext letters. Starting with column 1 for example, it can be seen from Figure 5 that letter "D" occurs 4.4 times (on the average) for every 100 letters in English, "F" occurs 2.8 times, "I" 7.4 times, and so on. I could have summed these up, but I learned in my studies that it is actually better to compute the sum of the logarithm of the frequencies instead. Thus, for the first column, $\log(4.4) + \log(4.4) + \log(3.8) + \log(7.4) + \dots = 7.3$. One advantage of using the log of the values rather than just the values themselves is that the overall sums will be penalized for containing letters B, J, K, Q, X and Z--that is, those with negative log frequency values--and will more dramatically reduce the overall score. The final log sums for each column are presented in Figure 12 (the table has been turned sideways for ease of presentation; the rows should correspond to the columns of Figure 11).

Again, the details of these calculations are not really important. The main point is that the rows in Figure 12 with the highest values are the most likely to correspond to the correct plaintext letters.

Figure 12. Log sum frequencies for each column permutation.

Words Start To Appear

It is easy to see from Figure 12 that the three highest log sum values correspond to the rows beginning with "R" (log sum=25.2), "T" (log sum=23.7), and "O" (log sum=24.8). These three values are significantly higher than any of the others. Now, I often imagine a professional cryptographer with a bag of special tools, not unlike a consummate locksmith or doctor, but consisting of alphabetic letter tables and charts rather than picks or scalpels. The English-language frequency chart and the I.C. were such tools, and now another one--the frequency table of initial letters common in English-language words--comes into play. In other words, how likely is it for a particular letter to begin a word? The letter "O" occurs as an initial letter an average of about 7.2 times per 100 words, and "R" only about 3.1 times. Because the letter "T" begins a word an average of 16 times per every 100 words, however, it is a good letter to try first.

Knowing that one of the most common words in English is the word "THE," and given that the first letter of Part II could very well be a "T," I naturally wanted to see what would happen if "H" were the second letter and "E" the third. But, once I chose these letters, I would then be committed to using all the other letters in those rows. The permutations with log sums for the second and third columns of Figure 8 are presented in Figures 13 and 14, respectively.

Figure 13. Probable plaintext for second column (red).

Figure 14. Probable plaintext for third column (red).

It was seen that the log sum frequency for the "H" row of Figure 13 is 29.4, and for the "E" row of Figure 14 is 30.5--certainly strong values, indicating that these rows were good candidates for being the correct plaintext letters. More importantly, when these three alphabets were lined up, no obvious impossible combinations were seen (Figure 15):

Figure 15. Probable plaintext for first three column of Part II.

What I was looking for here were contradictions, that is, combinations of letters that would not be possible in English, such as "AAA" or a "Q" that was not followed by the mandatory "U," for instance. Much to my delight, there were no contradictions found, and there were even heartening signs showing I was definitely on the right track. The word "AND" appeared in row 8, and an indication of the suffix "-TION" started to peek out from row 13. Best of all, that "KNO" in the 15th row was an unexpected gift, because I knew the odds were extremely good that there must be a "W" following it. All I needed to do to find the next plaintext column was to write out the set of permutations for the fourth column of Figure 8, find the set that contained a "W" in the 15th position, and insert this column into Figure 15.

Now that I was looking for logical words, calculating the log sum frequencies was no longer necessary. I just had to pick out the one row in each set of permutations that best fit with the already selected columns of Figure 15. What a treat not to have to add up any more of those log frequencies! (I had done so many of these that, at one point, I had memorized all the letter frequencies and proudly showed off this "skill" to my very patient family, friends and co-workers.) Continuing on in this manner, finding the remaining letters for columns 5-8 was not difficult, and the final plaintext of Part II quickly fell into place (the message is read horizontally):

Figure 16. Plaintext for all eight columns of Part II.

I noticed a few surprises. First, the letter "X" seemed to be serving as a punctuation mark in the fourth and 13th rows--probably a period. That could have confused the decryption if it had occurred in the first

few columns, but luckily it didn't. Second, I saw that the first word was not "THE" after all, but "THEY." This was another bit of luck (or serendipity), because my initial assumption was wrong, but it still led to the right answer. Finally, I noticed that the spelling of "underground" in the 10th row was "incorrect." I was so disconcerted by this that I made one of my rare trips (only about my third) out to the Kryptos sculpture to doublecheck my copy. I actually put my hand out to feel the letter and confirm that I had it correct. But I hesitated to call it an error--it could also have been a purposeful effort by the code's authors to make deciphering more difficult or possibly to provide a clue for interpreting a deeper part of the solution.

The next thing I did was to look at Part III of the code to see if it was encoded in the same way as Part II. Surprisingly, I quickly found readable plaintext forming from Part III, showing that what I was calling Parts II and III were really a single portion of code with a single encryption scheme. Although I wasn't expecting this to be the case, it was a pleasant shock to get the remainder of the message with almost no additional work.

When I tried to decode Part I using the same scheme, however, I found that it did not work. The first two lines of Part I were encrypted differently than the rest of the top section, so I had to start all over again, first computing the I.C.s under the assumption (which proved correct) that I was once more working with a Vigenere cipher for this part (Figure 17):

Figure 17. I.C. vs. number of alphabets for Part I.

Although the graph was cruder than for Part II because it was obtained from a fewer number of letters (note that some of the I.C.s seem to exceed the theoretical limit of 0.063), it was still obvious that a key length either 5 or 10 was being used. (These multiple high I.C. values were the result of the way the I.C.s were calculated. Because I.C. values are calculated by counting letter repetitions within a column, columns split up into multiples of the key length (5, 10, 15, and so forth) will show up as additional peaks at these values.) I tried both, but it was 10 that proved to crack this part. I performed the same analysis as just shown previously for Parts II and III, and the final decryption for the entire top half of the code is presented here (Figure 18):

Figure 18. Completed plaintext for top section of Kryptos code.

Deciphering the Bottom Section

Of course, all the time that I was working away at the top section of the Kryptos code, I would take occasional breaks to look at the bottom section. I saw almost immediately that this section, like the top, was made up of different parts. From the frequency by row charts (Figure 3, for example), there was a strong hint that the last three lines or so were encoded with a different scheme. So I decided to divide the bottom section into two parts at the question mark, and called them Parts IV and V.

Figure 19. Encoded bottom section of Kryptos code.

Starting with Part IV, the first thing I did was call out my old friend, the frequency count:

Figure 20. Frequency count for Part IV.

Comparing this chart to Figure 5, it was obvious that the letters in Part IV were occurring with the same frequency as they occur in normal English. This implied that this part of the code must be encoded with a transposition code (see Box A). Interesting. The Kryptos code (so far) seemed to be encoded with the two fundamental types of codes--the substitution and the transposition.

Because Part IV was almost certainly a transposition code, I knew that all 336 letters were already in plaintext, but I needed to determine their correct ordering. At first, it seemed like an impossible task, with a tremendously large number of possibilities. After reading up on the subject, however, I soon realized that there were standard ways to encrypt a transposition code, just as there were for the

substitution codes. And I had a hunch that the code's authors may have used one of the most fundamental types of encryption.

In a basic columnar substitution code, the plaintext message usually is first written out vertically in matrix form, with a certain number of columns and rows; it may or may not be a square matrix (see Figure 21). The columns are then mixed around, and the new ciphertext is read horizontally from the rows. To decipher the message, the reverse procedure has to be carried out.

Figure 21. Example of a basic transposition code.

Defining the Problem

To read the Kryptos transpositional message, then, the first thing that I needed to do was to determine the basic structure of the array; that is, ascertain into how many columns and rows the 336 letters were divided. For example, one solution might be to try 18 columns by 36 rows, because 18 times 36 equals 336. Another solution could be 56 columns by six rows. But the array does not need to be completely filled in; it could be, for example, a 17 x 20 grid, but with only 13 letters in the final row. So how would I know the right choice?

The big breakthrough took place when I realized that there were really only a finite number of possible solutions. I thought about the problem this way: suppose I constructed a simpler transposition message of, say, only 29 characters. (I arbitrarily chose 29 characters only as an illustration.) By investigating a much smaller message, I could concentrate on the mechanics of the code rather than be overwhelmed by the large number of characters involved. Further, because I had constructed this problem myself and therefore knew the answer, it would be possible to see the methodology required to solve it by working the problem backwards.

I imagined how things would look if I only had the ciphertext, without any clues on how to solve it (Figure 22). I would first need to figure out how many rows and columns there should be, decide how to break up the original message into "pieces" (columns) of a certain length (call this length "B"), and then figure out how many letters should be in each piece. Once this was known, I would have to assemble the columns into the correct order before I could read the plaintext.

Figure 22. Deciphering a simple transposition code with array not completely filled in.

Looking at Figure 22, it is apparent that, when the array is not completely filled in, there can only be pieces of length B or B-1, because the shorter pieces will always be just one letter shorter than the longer ones needed to fill in the array. Gradually, I realized that the mathematical relation should be as follows:

$$(B-1)x + By = 29,$$

or

$$y = (29 - Bx + x) / B$$

where x is the number of pieces of length B-1, and y is the number of pieces of length B. Because x and y have to be integers, (pieces cannot consist of fractions of a letter), the condition that there are only a finite number of solutions is forced. So, for a given B, only certain y values are possible for each x. For example, if B=5:

x	y
0	5.80
1	5.00
2	4.20
3	3.40
4	2.60
5	1.80
6	1.00
7	0.20
.	.
.	.
.	.

Because values of x greater than 7 will always yield y values that are less than 1.0 (and therefore not positive integers), in this example there are only two possible choices for x : 1 and 6.

I was ready to try my model on the actual Kryptos code. Substituting the number of characters in Part IV into the previous equation, I obtained the formula:

$$y = (336 - Bx + x) / B$$

I could then determine how many pieces of length B were possible to fit into the 336 characters inscribed on this part of the sculpture. I had noted that Part IV began with the word "END." This could be only a coincidence, but I wondered if it could be a hint by the code's author that maybe this word represented the final column in the grid. If true, this would imply either a grid made up mostly of columns of length 3 and 4 (that is, $B=4$), or else a grid made up all of column length 3 ($B=3$). The process of discovery is often intuitive as well as analytic, and, from working with the code for so long, I had gradually developed the strong feeling that the matrix seemed to consist of only a few long rows. My feelings were based on the realization that once the columns were finally assembled, they would need to be reordered into columns to make horizontal words. When I tried possible solutions using long columns, I found so many contradictions that I was forced to conclude that the columns could not be very long.

Finding the Columns

After many months of trying different ideas, based on the previous arguments, I finally decided to try $B=4$ as a working hypothesis. But I still needed to know how many words of length 4 and how many of length 3 there were. From symmetry, I had the feeling that the final three letters in Part IV represented a three-letter word, because I had hypothesized that the first three letters did. From the previous formula, I knew that there were only certain solutions that were allowable ($x=0, y=84$), ($x=4, y=81$), ($x=8, y=78$)...all the way up to ($x=112, y=0$). That is still 28 different solutions to check, but fortunately I tried them in order so that the correct choice of ($x=8, y=78$) was the third one that I considered:

Figure 23. Correct (but unordered) column arrangement for Part IV.

Re-ordering the Columns

I then needed to arrange the columns back into the correct order. (By the way, I should mention, if it is not already obvious, that I am presenting these steps in a certain order only for clarity. In reality, they were performed haphazardly, and I would often jump back and forth as new ideas came to me.) For this purpose, I reached into my cryptographer's bag of tricks one last time and used something called a

"digraphic frequency table." This chart depicts how often, on the average, any two alphabetic letters occur next to each other. For the first two columns in Figure 23, for example, the combination "YO" occurs 0.64 times per 1,000 letter pairs, "AH" occurs 0.13 times, and so on. By summing up these frequencies for each pair of columns, I was able to calculate which columns were most likely to occur next to each other and fit them back together like a jigsaw puzzle. The final solution is presented in Figures 24 and 25.

Figure 24. Ordered column for Part IV.

Figure 25. Correctly ordered text for Part IV.

The text appears to be a quote from a book by Howard Carter on the opening of King Tut's tomb. Note the troublesome "Q" at the end--I tripped over this more times than I can remember. When I first looked at this part of the code, I thought I could use the "Q" to help me solve the message, because I believed it would absolutely have to be followed by a "U." But this turned out to be a false lead (one of many).

Wrapping Things Up

Even though it was not necessary to determine the keywords used in the substitution code in order to read Parts I, II and III, I was curious to find out what they were. It was easy enough at this point: find the plaintext letters in the top row of the Vigenere Tableau in Figure 1, follow those columns down until the ciphertext letter is found, and then read off the first letters of those rows for the keyword letters. (Because the Kryptos cipher actually uses a modified version of the Vigenere code, the alphabets down the left-hand side and along the top of Figure 1 first have to be removed for this to work properly.) It turned out that the eight-letter keyword for the first two lines of Part I was "ABSCISSA," and the 10-letter keyword for Parts II and III was "PALIMPSEST." I knew what an abscissa was, but this other word was not familiar to me. I looked it up and found that it was "A written document...that has been written upon several times, often with remnants of earlier...writing still visible." Very interesting. Is the cipher trying to tell us that the Kryptos sculpture contains multiple layers of code--written one on top of the other?

I went back to the I.C. calculations for Part I and found a surprise when they were extended out to 16 alphabets:

Figure 26. I.C. vs. number of alphabets for Part I, extended to 16 alphabets.

I knew that the peak at $m=10$ was caused by the 10-letter keyword, and the peaks at 5 and 15 are expected artifacts of the keyword (at one-half and one-and-a-half times the keyword length, respectively). But what was this peak at $m=13$? Whatever it was, it also showed up in the I.C. plot for Part II (Figure 7). Could this be a sign of another code overlaid on top of the others--possibly one using a keyword of 13 letters? Or just a coincidence?

And what about Part V of the puzzle? I haven't spent much time on this portion [redacted] but there are interesting patterns in here as well. Although this part is no doubt more difficult than the previous pieces, I am confident that it is not impossible. I doubt that the code's authors would get much pleasure out of writing an unbreakable code.

Final Thoughts

I hope I have inspired some people to study the Kryptos puzzle and to give it a try. Even the parts of the code that already have been decrypted still have to be interpreted for their deeper meaning. There are many pieces to be put together and many layers to be peeled away.

Nothing is more frustrating than reading a cryptography book where the author easily and in a straightforward manner shows how a code was solved. It is a little like reading one of those books on "How I made a million dollars." The methods typically work for that particular set of circumstances, but they often do not work in your particular case. Similarly, codes have distinct characteristics that

frequently require each to be solved with a unique method. Because I did not see any need to dwell on the hundreds of things I had tried that didn't work, the methodologies I presented may seem a little too straightforward and deceptively easy. As the adage goes, "If at first you don't succeed"--and in cryptology you never do--"try, try again." If the methods that I tried had not worked, I would have tried others until I found something that did.

I genuinely enjoyed working on the Kryptos ciphers. Professional cryptographers almost certainly could have broken these codes much faster, and would have used superior methods. But I doubt that they would have derived as much satisfaction as I have. I didn't use any computers to decrypt the Kryptos codes--just pencil and paper, some common sense, and a lot of perseverance. Using a computer would have cheated me out of the feeling of accomplishment that I obtained, because I've found that often in life the journey itself can be more gratifying than arriving at the final destination. Mountains are not climbed nor marathons run merely to reach a geographical location--there are much easier ways to accomplish these feats--but as personal and spiritual challenges to the participants.

When confronted with a puzzle or problem, we sometimes can lose sight of the fact that we have issued a challenge to ourselves--not to our tools. And before we automatically reach for our computers, we sometimes need to remember that we already possess the most essential and powerful problem-solving tool within our own minds.

is in the Directorate of Intelligence.

Unclassified

[Next](#) [Previous](#) [Contents](#)