At Bletchley Park

WILLIAM T. TUTTE

August 16, 2000

I Introduction

It was January of 1941. Britain was at war with Germany and Italy. Germany had conquered Western Europe but Italy was suffering defeat in Greece, Egypt and Libya. I was a research student at Trinity College, Cambridge. I remember an interview with my tutor, Patrick Duff, when he told me that I should go to a certain town about 50 miles away for an interview about a possible war job.

Gaps occur in one's memory after 50 years or so. I must have had the interview and something of interest must have been said in it. But I remember nothing of it. I fancy that town must have been Bletchley, but feel no surety even of that. However I soon found myself in a cryptographical school in London, at a place near St. James' Park and his Underground station. There I learned how to deal with cryptograms of the First World War.

Why did they pick me? I think it was because I had acquired a reputation as a solver of puzzles at Trinity. I had been one of the four undergraduates who had, for fun, worked out a solution of the problem of dissecting a square into unequal squares. Curiously R. Sprague in Berlin had worked out a solution at about the same time. (His square was the first to be published).

Presumably my performance at the school was found satisfactory and I was recruited to Bletchley Park with the rank of Temporary Assistant Clerk (of the Foreign Office). I was a member of a small group called the Research Section, which had a room in the

Mansion. But after a year or so it moved into one of the temporary huts. I do remember the names and appearance of at least some of my fellow-members. There was Captain Gerry Morgan, the Head of the Section and his brother Stanley, a lieutenant. There was Sergeant Rylands, smoking his pipe. There was Daphne Bradshaw whose husband was also at the Park as an administrator. Others I remember but think they joined later. I noticed that the differences of Army rank seemed of no importance within the Section.

Bletchley Park housed an organization whose work was the study of enemy codes and ciphers, with the object of reading them. Some of the people there were members of the armed forces. Others, like myself were civilians. We mixed together in the Research Section on equal terms. It is well-known that cryptographers at Bletchley were very successful with a German machine-cipher called ENIGMA. There was another machine-cipher code-named FISH. The traffic on its first radio-link to be studied was called TUNNY and other fish-names were given to links that came later. In a book called Codebreakers, written by Bletchleyites, it is explained that ENIGMA and FISH were complementary from an Allied point of view. Navy and Air Force ENIGMA yielded satisfactorily to the cryptanalysts but Army ENIGMA was recalcitrant. However FISH was exclusively for the Army, and it was broken. It is also explained that FISH was a high-level cipher, typically between generals in the field and their superiors in Berlin. In the later years of the War the only available methods of breaking FISH involved too much computation to be done entirely by hand, and so the computation had to be done mechanically. The mechanization culminated in the pioneering electronic computer called COLOSSUS.

I should like to emphasize here that COLOSSUS worked on FISH, not on ENIGMA as is sometimes stated. ENIGMA had its own helpful machines, but they were called "bombes". One reason why the two groups of cipher-breakers needed different machines was alphabetical; ENIGMA used the ordinary German alphabet and FISH used the teleprinter alphabet of 32 letters.

It is remarked somewhere in Codebreakers that the Germans referred to one of their radio links as Sägerfisch, and that Bletchley's term FISH derived from this.

The Research Section was consulted by other Sections about ciphers that were still unbroken or whose exploitation had not yet been reduced to a routine. When I arrived

the Section was chiefly concerned with an Italian naval cipher of the latter kind. The encipherment was done on a commercial machine of Swedish manufacture. We called it a Hagelin machine. I took it that a Mr. Hagelin was either the inventor or the manufacturer. There being nothing secret about the machine, specimens were available at Bletchley.

II The Hagelin machine

The Hagelin machine produced what we called an "additive cipher". It used a 26-letter alphabet and the letters were taken as equivalent to the integers from 1 to 26 in their numerical order. With this convention letters could be added modulo 26.

The mechanical details have become vague in my memory. But somehow the operator fed plain text into the machine. The machine generated a sequence of letters, of random appearance, that we called the key. Letter by letter the machine added the key to the plain text and the result was the cipher text. This would be sent by radio from originator to recipient. The latter would set his machine to subtract the key from the cipher text, and so recover the plain text. If the message was intercepted by a British station, say at Malta, the cipher text would go to the Research Section at Bletchley Park with, of course, no information about the key.

There were many possible initial settings of the machine, and each one gave a different key. The originator could choose one of these settings and then he had to tell the recipient which one he was using. The sensible method would be to say in clear "I am using setting number n" and then the recipient would look up setting number n in a secret book. I suppose this method was used but cannot remember. Perhaps it was something less secure.

The great weakness of an additive cipher is the danger that two messages may be sent on the same setting, i.e. same key. Cryptographers call this a depth of 2. Three of the same setting would be a depth of three, and so on. Let us denote the common key of two messages in depth by K. Let the plain texts be P_1 and P_2, yielding cipher texts C_1 and C_2 respectively. Then we have two equations

$$P_1 + K = C_1 \qquad \text{and} \qquad (1)$$

$$P_2 + K = C_2, \tag{2}$$

the addition being mod 26. Subtracting the second equation from the first we get

$$P_1 - P_2 = C_1 - C_2 \tag{3}$$

So the cryptanalyst only has to subtract one cipher text from the other to get the difference of the two plain texts. The machine itself has vanished from his problem: all he has to do is disentangle the difference $P_1 - P_2$ into the two plain texts P_1 and P_2. How did he know the two messages were on the same key? The originator had to say so in his preliminary clear messages to the recipient.

Disentangling P_1 and P_2 is simple in principle if not always easy in practice. Each clear message usually had an address very near the beginning. A likely one for a message going to Italy would be KSUPERMARINAKALTK. The Italian operators, having no other use for the letter K, made it a word spacer. One would put such an address into P_1 and calculate the corresponding part of P_2 from one's knowledge of $P_1 - P_2$. If that part made Italian naval sense, then that section of P_1 and P_2 could be assumed known. The next step was to guess a continuation of P_2 and verify it (or otherwise) by checking against P_1 and so on. With luck P_1 and P_2 could be read up to the point where one of them ended, and the information could be sent to those who would decide how to use it.

The cryptanalyst would be more interested in the fact that, having P_1 and C_1 he could now get the key K from Equation (1). He could then contemplate the key and the machine and brood over the problem of what configuration of the machine would produce that key.

Users of additive cipher machines are, or ought to be, strictly forbidden to send two messages in depth. But evidence from the Second World War suggests that perfect obedience is difficult to get in practice. I confess I do not remember how good the Hagelin users were at avoiding depths. I do not remember having to break one myself but the danger was there and my explanation of it will, I hope, be helpful when I come to the discussion of another additive cipher machine.

Let me now write about the structure of the key. First it was the sum of six periodic subkeys, which I will denote by $K_1, K_2, \ldots, K_6$. Each of these had a period of the order of 30. Each key K_i had an associated small positive integer n_i and was simply a sequence of numbers each of which was either 0 or n_i. We called these subkeys "wheels". For each of them was generated by an actual mechanical wheel with pins that could be punched in or out. The position in or out of such a pin would determine whether K_i had 0 or n_i in that position. It was also possible to adjust n_i for each K_i. Normally the sum of the n_i would not exceed 26. So the addition of the six subkeys, or wheels, could be regarded as ordinary addition, not mod 26. Also the n_i went in non-decreasing order from K_1 to K_6.

The number n_i and the arrangement of pins made the "wheel-pattern" of K_i, its periodic sequence. Wheel-patterns would change from month to month, or from day to day, according as to how security-conscious were the authorities, but not from message to message.

My task with the Hagelin machine came after the wheel-patterns had been determined, I suppose by analyzing a key derived from a depth. I was given messages to set on the known wheel patterns. That is, I had to find which pin on each wheel was active in the initial setting. I would guess an address in a likely part of the message and calculate, on that assumption, the corresponding part of the key. Some numbers in this key would be so high as to require a contribution of n_6 from K_6. Others would be so low as to demand 0 from K_6. So with luck K_6 could be set or the assumption could be disproved. Having succeeded with K_6, I would go on to treat K_5 similarly, and so on. I became fairly proficient at this. Sometimes, that is some nights, I would be invited to sleep in the Italian Naval Hut so that I could be awakened and consulted over some particularly awkward message.

It came to pass that the breaking of these messages was reduced to routine, or so nearly so that they needed no longer to be the concern of the Research Section.

III I meet TUNNY

I think it was in October 1941 that I was introduced to a new German cipher called "TUNNY", the first of the "FISH" links. The intercepted cipher messages were sequences of letters of the International Teleprinter Alphabet. A letter of this alphabet could be represented conveniently as a column of five symbols, each restricted to two values, say 0 and 1. The letters could be described mathematically as 5-vectors mod 2 and could be added as such. Thus G and C in the teleprinter alphabet are

$$\begin{Bmatrix} 0 \\ 1 \\ 0 \\ 1 \\ 1 \end{Bmatrix} \text{ and } \begin{Bmatrix} 0 \\ 1 \\ 1 \\ 1 \\ 0 \end{Bmatrix}, \text{ respectively,}$$

and with the above convention their sum is

$$\begin{Bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 1 \end{Bmatrix}$$

which is the letter H.

The vectors can of course be written as rows instead of columns. The symbols 0 and 1 appearing in a letter, taken in order down the column, or from left to right along the row, were called at Bletchley its first, second, third, fourth and fifth "impulses". In a sequence of letters such as the plain or cipher text the corresponding sequence of n^{th} impulses of the letters was called the n^{th} "impulse" of the letter-sequence. $(n = 1, 2, 3, 4, 5)$.

For some reason cryptographers at Bletchley preferred to write 1 and 0 as cross and dot ($\times$ and $\cdot$) respectively. Electrical engineers used -1 and 1 and multiplied in the ordinary way where we added mod 2. (It came to the same thing.) The operation we

called adding was easily done electrically and so was quite likely to be used in a teleprinter cipher. The 32 letters of the alphabet in row form are as follows:

· · · · ·	All-space	× · · · ·	E
· · · · ×	T	× · · · ×	Z
· · · × ·	Carriage Return	× · × · ·	D
· · · × ×	O	× · · × ×	B
· · × · ·	Word space	× · × · ·	S
· · × · ×	H	× · × · ×	Y
· · × × ·	N	× · × × ·	F
· · × × ×	M	× · × × ×	X
· × · · ·	Line feed	× × · · ·	A
· × · · ×	L	× × · · ×	W
· × · × ·	R	× × · × ·	J
· × · × ×	G	× × · × ×	Figures
· × × · ·	I	× × × · ·	U
· × × · ×	P	× × × · ×	Q
· × × × ·	C	× × × × ·	K
· × × × ×	V	× × × × ×	Letters

Some of these "letters" are instructions to the machine. The one called "Figures" (× × · × ×) tells the recipient to read the following letters as digits according to the standard code. The one called "Letters" (× × × × ×) tells him to stop this and go back to letters.

Each cipher message of TUNNY was preceded by a sequence of twelve letters belonging to the ordinary German Alphabet. As a safeguard against error these might be expanded into common personal names. (A for Anton, B for Bertha, C for Caesar, D for Dora, E for Emil, F for Friedrich and so on). This sequence of twelve letters we called the "indicator". Presumably it told the recipient how to set his machine.

Sometimes two messages would be found to have the same indicator. Some of these pairs were successfully read on the assumption that the cipher was additive and used the

mod 2 addition described above. It could now be asserted that TUNNY was an additive cipher and that the twelve indicators determined the key. Very likely they specified the settings of twelve wheels, something like Hagelin wheels.

If θ is a letter of the teleprinter alphabet it satisfies the equation

$$\theta + \theta = 0 \text{ (i.e. (} \cdot \ \cdot \ \cdot \ \cdot \ \cdot \text{), All-space).}$$

Thus each letter is its own negative; subtraction is the same operation as addition. This fact introduces new difficulties into the process of reading depths. Equation (3) is now replaced by

$$P_1 + P_2 = C_1 + C_2 \tag{4}$$

The equation has acquired more symmetry. Now even if we have found P_1 and P_2 the equation does not tell us which plain text goes with which cipher text. (P_2 may be the true P_1). However if the messages are read up to the point where one cipher text ends, then the plain text that obviously ends there must go with the shorter cipher text. So in that case the ambiguity may be resolved.

But in practice a depth is likely to be read in patches, with short or long gaps between consecutive patches. Perhaps a gap represents letters that, because of some natural interference, were not received at the British interception station, or were received only in a damaged form. Perhaps it appeared only because depth-reading is difficult. Then the ambiguity applies to each patch separately. There is nothing in the algebra to show which part-message of one patch goes with which part-message of another. Of course the sense of the part-messages, if they are long enough, may indicate which goes with which.

On one occasion a German operator made a major error. A message about 4000 letters long was sent. There must have been some difficulty of reception, for the message was sent again. The sender made the mistake of using the same indicator as before. If the second version had been an exact copy of the first this would not have mattered. But evidently the sender did not insert a tape; he typed his message in letter by letter. His punctuation was different. He put in fewer word-spaces. So the old version was drawn out to a greater length than the new. The two messages were vulnerable to a cryptanalyst as a depth of two. Moreover once they were recognized as differently punctuated versions

of the same text the reading of the depth was made much easier. Message 1 was about to say what Message 2 had said a little way back. The depth was read by Colonel John Tiltman, next above Gerry Morgan in Bletchley's hierarchy. He even decided which plain text went with each cipher text; the one that had the complete message went with the shorter cipher text. Adding P to C he got a piece of TUNNY key four thousand letters long.

All this happened two or three months before I met TUNNY. In those few months other Bletchleyites were trying to "break the TUNNY key", that is, to describe mathematically a machine that could produce that stretch of key. They worked hard, guided by an ingenious theory, but to no avail. Some time later I was told about this theory, but it is now gone from my memory. By then it was known to be wrong.

Was it a gesture of despair that Captain Morgan, that day in October, handed me the TUNNY key, with associated documents and said "See what you can do with this"? Apart from the 4000 letters of key there was just one other item of information that struck me as possibly significant. It was that 11 letters of the indicator could take all 25 values permitted by the German alphabet, (omitting J). But the other letter took only 23 values. Did that mean that there was a wheel of period 23?

In the cryptographical school in London we often attacked our simulated cipher messages by writing them out on some period that seemed appropriate. Thus if the period was 29 we would write the first 29 letters in a row on squared paper, then the next 29 exactly below them, and so on. We would then look for repetitions from row to row. If there were significantly many we knew we were well on the way to a solution. So I thought I would start on the TUNNY key by writing it out on a period, or perhaps only one impulse of it. I do not think I had much faith in this procedure but I thought it best to seem busy.

I chose to write the first impulse on some period. It would be less trouble to write a single impulse than to write the complete key. Besides, just possibly the Part might be cryptographically simpler than the Whole. But what period should I use? That information about the indicators suggested that 23 might be worth trying. So, but very doubtfully, might 25. Why not try both at once? 23×25 was only 575 and there were

4000 letters of key. So I wrote the first impulse of the key on a period of 575, filling about 7 rows.

The next step was to look for repetitions of sequences of five or more dot-cross symbols from row to row, the second group of the repetition exactly under the first. But there were not significantly many of these repetitions. Before taking the next logical step, that is, doing likewise with the second impulse, I checked the diagonals. It seemed that I would have got good results on a period of 574.

I wrote the first impulse on a period of 574 and marveled at the many repetitions down the columns from row to row. But surely the Germans would not use a wheel of that length? Perhaps the true period was 41, this being a prime factor of 574? So I wrote the first impulse a third time, now on a period of 41. I got a rectangle of dots and crosses that was replete with repetitions.

My interpretation of this effect was simple. The first impulse, K_1, was the sum of two sequences of dots and crosses. The first was periodic, of period 41. The second had the curious property that in it dot was more likely to be followed by dot than by cross, and similarly cross was more likely to be followed by cross than by dot. Careful study showed that the second sequence could be interpreted as the product of a wheel that sometimes moved on from one letter of text to the next and sometimes stayed still. The first sequence of course came from a wheel that always moved on one place from one letter to the next. I called the first wheel, of period 41, the χ-wheel and the second, which turned out to have a period 43, the ψ-wheel. These names stuck.

At this stage the rest of the Research Section joined in the attack. Each other impulse of K was split into its χ and ψ components. It was found that the five ψ-wheels moved in step. From one letter to the next either they all moved on one place or they all stayed still. Besides the χ- and ψ-wheels there were two "motor wheels". One of period 61 moved regularly like the χ-wheels. The other moved on one place when the first showed cross and stayed still when the first showed dot. This second motor wheel had period 37. The ψ-wheels moved on when it showed cross and stayed still when it showed dot.

Thus were the entire workings of the TUNNY machine exposed without any actual physical machine or manual thereof coming into our hands.

The reader will observe that the Germans could have arranged the patterns of the ψ-wheels so that in the resulting ψ-keys consecutive symbols were as likely to be different as similar. Then my method would have failed. We must class their failure to do so as their second major mistake, the first being the long message repeated with the same key. Either mistake alone I think they would have got away with. But the two together proved fatal.

About that indicator position that allowed only 23 letters. There was a corresponding wheel of period 23. It was the fifth χ-wheel. If I had not noticed those diagonal repeats and had accordingly applied my method to the other impulses it would have succeeded on the fifth. Then I suppose my success would have been attributed entirely to close logical reasoning. As things were, I was supposed to have had a stroke of undeserved luck. (Think twice, O Gentle Reader, before thou takest an unexpected and opportune short cut.)

IV The days of the indicator

The indicator of that long depth was HQIBPEXEZMUG. Now that we knew the workings of the machine, albeit only in mathematical abstraction, we could analyze keys obtained from shorter depths. This involved writing each impulse of the key on the period of its χ-wheel, now known, and inferring for each consecutive pair of χ-symbols whether they were alike or different. If we assumed, and were right in our assumption, that the new key had the same wheel-patterns as HQIBPEXEZMUG, the problem became one of finding not wheel-patterns but "wheel-settings". We had to find for each wheel what symbol of its periodic pattern corresponded to the initial letter of the message. From a few successes of this kind we found that the wheel-patterns changed only from month to month. The wheel-settings of course changed from message to message, save when an operator carelessly sent a depth.

Our decodes were of messages too old to be of interest to our customers. There was a little current traffic but it provided no depths. Two or three months into 1942 there was more traffic and our German suppliers sent a few depths. There came one that seemed

usable, being about 1000 letters long.

The depth was broken and the five impulses of the key could now be written on the appropriate periods. Alas, the resulting rectangles were just random-seeming arrays of dots and crosses. We supposed that the Germans had noticed the weakness of their early ψ-patterns and had corrected it. If so their ψ-patterns now had more changes than continuations, just sufficiently more to make my method of analysis ineffective. The alternative hypothesis that they had switched to an entirely different machine was too awful to contemplate.

I was now immersed in the theory of TUNNY and, I suppose, worried as much as anyone in the Section about this new impasse. Looking down the list of indicators one day, I noticed two that differed only in one of the twelve letters. This was something new, a "near-depth". I wondered if a near-depth could be read like a real depth, though no doubt with much greater difficulty.

If the old rules were still valid the two keys differed only in the setting of one χ-wheel. I do not now remember which one. For the purposes of explanation let me pretend it was the first, with period 41. Then the difference Δ between the two keys would be the mod 2 sum of the two versions of χ_1, corresponding to two different wheel-settings. It would therefore have period 41. Once 41 places of Δ had been determined, the near-depth could be made effectively a true depth by adding Δ to one of the messages. A successful reading would give not only many letters of key but extra information about the first χ-wheel. I was optimistic that this extra information would make possible the breaking of the key.

So it turned out. We, being incapable of reading near-depths, asked for help from another Section, one that abounded in linguists instead of mathematicians. At first the linguists were reluctant to tackle the problem. This was understandable: if you guessed a letter in one message there were now two alternatives for the corresponding letter in the other. But they set to and succeeded. After reading a stretch of 41 (or whatever it was) letters they converted the pair of messages to a true depth and carried on, with much reduced difficulty, to the end.

The problem now was what to do with Δ. I will go on pretending that Δ belonged to the first impulse. Then it gave rise to 20 possible patterns for χ_1, one for each of the

possible relative settings of the first χ-wheel in the two messages (irrespective of which setting went with which message). At first sight it seemed that the 20 should be 40 since the χ-pattern could be reversed without altering Δ. (To reverse a pattern, change dot to cross and cross to dot). However it would not worry us if we were using not the German χ-pattern but its reverse. We would then use the reversed ψ-pattern of that impulse too, and the key would be correct. So there were 20 really different possibilities.

Most of the possibilities could be rejected as implausible. We were sure that a χ-pattern should have, roughly, as many crosses as dots and as many continuations as changes. Judging by the χ-patterns we had found for the old messages, the German operators agreed with us. We also imagined that the composers of wheel-patterns received some such advice as this: "Remember, Conrad, no excessively long blocks of dots, or of crosses." So the 20 possibilities reduced to very few plausible ones. For each of these we had the χ-pattern and its settings in the two messages. But because of the peculiarities of mod 2 addition we did not know which setting went with which message. We would have to try both possibilities.

Suppose we are lucky enough to apply the correct χ_1-pattern in its correct setting to the first impulse of key. By addition we get the first ψ-impulse, ψ_1. When this has two different consecutive symbols we know that the ψ-wheels have moved from one symbol to the next. But when ψ_1 has two consecutive symbols the same we can say that, more likely than not, the ψ-wheels have stayed still from one symbol to the next. This is because we believe that the Germans have now put more changes than continuations into their basic ψ-patterns.

Since the five ψ-wheels move in step we can apply the knowledge of ψ-movement thus obtained, to any other impulse. Let us say the second, in which the χ-period is 31. For each consecutive pair of symbols there we can say either "Here the ψ-wheels have certainly moved" or "Here, more likely than not, they have stayed still." In the first case ψ_2 most likely shows a change, and in the second a repetition. This is because of the new German rules of ψ. Using this information we can deduce, pretty accurately, the succession of changes and repetitions in the χ_2-pattern. This in turn, enables us to correct that succession in the ψ_2 sequence, making it more accurate if not perfectly so.

Then we can go to the third impulse and get very accurate information about χ_3, and so on to the breaking of the key.

If there is more than one plausible pattern for χ_1 we simply have to try each possibility in turn, until the key collapses.

By such a procedure we broke the key of the near-depth, getting the wheel-patterns of the corresponding month. I remember that a later near-depth gave us the patterns for a later month.

By this time the Research Section had moved from a room in the Mansion to one of the huts. I do not remember the number but I was told recently at Bletchley that it must have been Hut 15.

At this stage Alan Turing, tiring perhaps of ENIGMA, made a few weeks visit to the Research Section. He became interested in the problem of breaking a true TUNNY depth and he found a method of doing so. Some people, though not myself, became expert in Turing's Method. It seemed to me that it was more artistic than mathematical. To start with, you picked a pair of consecutive positions in some χ-wheel and assumed either a change or a repetition there. Whichever you chose you had a 50 per cent chance of being right. The pair would recur, with interval the χ-period, throughout the key. Let us suppose χ_5 chosen so as to have the shortest period, 23, and the greatest number of recurrences. At each recurrence you could calculate the corresponding pair in ψ_5 and say whether it most likely represented a ψ-motion or a ψ-stoppage. Since the ψ-wheels moved in step this information applied to all the impulses. You assumed the "most likely" alternative for each recurrence. Then you could deduce information about some consecutive pair in the other χ-patterns and extend it through the key on their periods. Soon of course you got clashes. At each clash you rejected the alternative you felt in your bones was the less likely. With sufficiently reliable bones you might break all the wheel-patterns. If you failed you tried again, perhaps with a different pair of symbols and perhaps with the other assumption about the first pair.

So by Turing's method and by the method of near-depths we got our monthly wheel-patterns and whetted the appetites of our customers with the messages of the depths and near-depths. But those customers wanted more deciphered messages than that. So we had

to work on another problem. Given a message not in depth, and given the wheel-patterns, find its wheel-settings.

I said in a previous essay that I could not remember how this problem was solved. Yet I do remember that I was sure it could be solved even before I found the near-depth. I remember saying to those linguists "If we can break this near-depth, we shall go on to break many more messages" But now a memory of Sergeant Rylands has revived. Referring to a still-famous poster of the First World War, he said "If ever that little boy asks me 'What did you do in the Great War, Daddy?' I shall have to say 'I dragged HERRING'" (or it may have been some other FISH-link). "Yes", I say, "We, and notably Sergeant Rylands, solved that problem by the device of dragging".

I should mention that the indicator letters were not in alphabetical order on their wheels.

We knew what indicator letters corresponded to the wheel-settings of any depth or near-depth we had broken. For the next break we would choose a message not in depth but with at least one χ-indicator for which the corresponding wheel-setting was known. I will now try to explain what could be done if the setting of one χ-wheel, say χ_1, was known. But I shall rely at least as much on reasoning in this year 2000 as on memories of 1942.

Messages of that time commonly started with an address. There being few addresses in use on a given link a cryptanalyst could use them as cribs. Of a given message he could say "This quite likely has 9OBERKOMMANDO9WEHRMACHT9 very near the beginning." It was customary with us to write the word-space symbol as 9. Adding this address to the message in a possible position we would get a stretch of possible key. On the assumption that it really was key we would add the known χ_1 to the first impulse and so get ψ_1. Then in every impulse we could say of each consecutive pair of symbols whether it most likely represented a ψ-movement or a ψ-stoppage. Then for each consecutive pair of symbols of χ_2 we could say whether they were most likely to be a change or a continuation. We would then slide the actual χ_2-pattern along χ_2 until it clicked, until "most likely" agreed sufficiently often with "actual". If there was clicking for the other three impulses as well, then the five χ-settings had been found. If there was no clicking,

then the address had to be "dragged" to the next position. If the settings of two χ-wheels happened to be known from the indicators then the process became much quicker.

When the χ-wheels had been set the χ-key could be added to the part of K under the address to get the corresponding part of the ψ-key. There the pattern of ψ-movement and ψ-stoppage would be known very accurately and so the ψ-wheels could be set too. After this the reading of the message could be completed. We would use the fact that there were just two possibilities for the next ψ-letter since the ψ-wheels either all moved on one place or all stayed still. The corresponding χ-letter would be known since χ-wheels moved regularly. So there would be two possibilities for the plain language letter and usually there would be no doubt which was right. When the entire ψ-key was known it would be possible to tackle the motor wheels.

With each new success the process of dragging became easier since the meanings of more indicator letters were known. If the crucial depth or near-depth came early in the month we might soon find ourselves reading current traffic.

Waiting for that depth could be exasperating. I began to dream of breaking wheel-patterns without any depth or near-depth, using only information from the indicators. All messages having the same letter in the same indicator place would have the same setting of the corresponding wheel. Perhaps, given enough messages of the same month, that information could break the wheel-patterns.

Consider the first letter, say in the third impulse. The first clear letter was likely to be word-space ($\cdot\ \cdot\ \times\ \cdot\ \cdot$). One would assume it to be enciphered by the first χ-symbol and the first ψ-symbol in relation to the indicator-letters concerned. Assuming the first χ_5 symbol at indicator A to be cross, we should be able to calculate the ψ_3-symbols at the appropriate indicator-letters by studying the set of messages with χ_3-letter A. From these we could expect to get all the first symbols at the χ_3-letters, and so on just to check. The process would of course be gone through for all the five impulses. Garbles and eccentric first-clear letters could be circumvented if few enough.

I remember that this process worked well for the first letter in the past month that I chose to study. It gave the unexpected information that the first letter of a message was enciphered by the χ-key only. This cleared up one of our minor problems. It explained

why the first letters of our decodes so often seemed to be wrong. It also made it possible to resolve a familiar ambiguity: did we have the German wheel-patterns in χ and ψ or their reverses? Not that anyone thought that resolution important.

The process could now be applied to the second symbols from each indicator-letter. It could now be assumed that the ψ-symbols could all be the same for a given impulse and indicator letter. I remember that this attack succeeded too. Now the comparatively few plain-text letters that were not word-spaces could be deciphered. If they could be interpreted as the first letters of addresses then some third clear letters could be guessed. Unfortunately the ambiguity between German and reverse reappeared with this second letter. Eventually it would have to be resolved, for correlation with the first χ-symbols. Meanwhile one had to accept that there were two alternatives.

Difficulties increased with the third symbol from the indicator letter. Some messages would be enciphered from the second symbol of the active ψ-key and others from the first. The letter of the plain text would be less likely to be word-space. But the other likely plain-text letters were few and could be guessed. There seemed to be good hope of deciphering the third letters. There was another comforting thought. No new independent ambiguity should appear. The third letter could be correlated with the second by way of the messages for which the ψ-wheels had stayed still.

My memory tells me that I coped reasonably well with the third letter but got confused with the fourth and made no further progress. Had that been the whole story it would hardly have been worth mentioning in this essay. But my attempt aroused the interest of a new member of the Research Section, James Wyllie. Wyllie was an editor of the Oxford Latin Dictionary. He was willing to take over from me, for he said "This is just the task for a lexicographer!" He too chose a past month and he carried the procedure so far that he was able to put his wheel-pieces together to make the complete wheel-patterns. So I have called this method of TUNNY-breaking "Wyllie's Method".

In 1996 when I first began to be seriously questioned about my memories of BP I was unable to say just how much practical use had been made of Wyllie's Method. But since then I have seen a report issued by the section called the Newmanry soon after the German collapse. I gathered that the traffic of several months had been broken, at least

partly, by that method. Its disadvantage was that it needed a large number of messages so it could only succeed late in a month. It might be cut short by the breaking of a depth or a near-depth. But even so its fragmentary wheel patterns would be of help in the process of dragging.

So the deciphering of FISH went on into 1943. Then there came a Black Day. From then on the cipher messages came to us without those helpful letters. They were replaced by a simple number. No doubt a German cipher officer would look up that number in a little book and find twelve letters printed against it. But we did not have that little book. It seemed that just one way of attack was still open to us. We could still recognize depths; messages in depth had the same number. Near-depths might still occur but we had no way of recognizing them. Production stopped, save for the occasional pair of messages in depth.

I suppose some German inspector had examined the process of encipherment and had exclaimed somewhat as follows. "Hey, you're giving those bastards information that you don't need to give them! I don't suppose it has done them any good, but it's wrong in principle. Stop sending the twelve letters!" Here I have assumed the inspector to have complete confidence in the security of FISH, believing that otherwise he would have demanded much more. On second thoughts it is easy to imagine that he did demand more but was overruled by his superiors.[1]

Be that as it may we had met with disaster. In spiritual metaphor we shouldered our pencils and squared paper and trudged glumly out of Eden.

In some such manner did that unknown German gentleman, as judged at BP, set going the Computer Revolution.

[1]After the German collapse an anecdote came to BP, I know not how reliable its source. A German mathematician had queried the security of a cipher machine early in the War. An army officer had replied "So what? We're winning the War, aren't we?" ("Famous Last Words", quipped Gerry Morgan.)

V The winter of our discontent

We still got, at rare intervals, depths that could be broken by Turing's Method. Then we would have the wheel-patterns for the current month. We asked "How can we set these wheels on the messages that are not in depth?"

Dragging without a known wheel-setting was not seriously considered. In theory I suppose 23 cryptographers could each have been given one of the 23 possible settings of the 5th χ-wheel and told to drag on the assumption that that was the correct one. That would have been an extravagant use of manpower, especially as the breaking of one message would no longer give wheel-settings for others.

The Research Station had now been joined by Professor Max Newman. Most of us worked together in one big room that occupied most of our Hut, but Gerry Morgan had his own enclosed office. There he and Max would foregather and discuss the problem. I gathered that they had an inkling of a possible method, which at times they would explain voluably to me and to others. I fear that I understood little of it and had little confidence in that little. For myself I dreamed of applying methods of linear algebra.

I have been asked "When did the study of $\Delta\chi$ and $\Delta\psi$ become usual at Bletchley?" Given a sequence S of teleprinter letters, the n^{th} letter of the sequence ΔS is defined as the sum (i.e., the difference) of the n^{th} and $(n+1)^{th}$ letters of S. Similarly can the sequence T of dots and crosses in a single impulse be converted into ΔT. Earlier in this essay I have been much concerned about whether two consecutive symbols in a stream T of dots and crosses were alike or different. Equivalently but more simply I could have asked whether ΔT had a dot or a cross. It seems that the Δ operation could have been used with advantage in Turing 's Method. I do not remember, but find it hard to doubt that it was. Using ΔK_1 instead of K_1 would certainly have helped in the breaking of the key of HQIBPEXEZMUG. Had I written ΔK_1 on a period of 41 each column would have been either predominantly dot or predominantly cross, and the dot-cross sequence would have been the pattern of the periodic sequence $\Delta\chi_1$. So I was probably doing no new thing when, in our emergency, I meditated on the possible use of Δ-sequences. But it is from that time that they are prominent in my memory.

A Δ-sequence of particular interest was $\Delta\psi$. It could be assumed that in encipherment the ψ-wheels would advance about half the time and stay still half the time. Then about half the letters in $\Delta\psi$ would be all-space ($\cdot\ \cdot\ \cdot\ \cdot\ \cdot$). Considering only the i^{th} impulse, $\Delta\psi_i$ would be dot when the ψ-wheels stayed still and it would sometimes be dot when they moved on.

I deltaed the encipherment equation for the first impulse of a message:

$$\Delta C_1 = \Delta P_1 + \Delta\chi_1 + \Delta\psi_1 \tag{5}$$

After contemplating this for a while without enlightenment I wondered if I could play one impulse against another since the ψ-wheels moved in step. So I wrote the corresponding equation for the second impulse

$$\Delta C_2 + \Delta P_2 + \Delta\chi_2 + \Delta\psi_2 \tag{6}$$

I noted that $\Delta\psi_1$ and $\Delta\psi_2$ were very much alike. When the ψ-wheels stayed still they were both dot. When the ψ-wheels moved on, wouldn't $\Delta\psi_1$ and $\Delta\psi_2$ be alike as often as not? I calculated that the sequence $\Delta\psi_1 + \Delta\psi_2 = \Delta(\psi_1 + \psi_2)$ would be about 70% dot. It seemed that I could achieve an imperfect elimination of ψ by adding my two equations together. Δ is distributive over addition so I could write my result as

$$\Delta(C_1 + C_2) = \Delta(P_1 + P_2) + \Delta(\chi_1 + \chi_2) + \Delta(\psi_1 + \psi_2). \tag{7}$$

All now, I thought, depended on $\Delta(P_1 + P_2)$. I investigated it for some known plain text and was delighted, but surprised, to find that it was as a rule about 60% dot. The upshot was that $\Delta(C_1 + C_2)$ and $\Delta(\chi_1 + \chi_2)$ would agree with one another more often than not. In favourable cases, such as one deciphered message that I checked, there might be 55% agreement.

Here was a method of wheel-setting! $\Delta(\chi_1 + \chi_2)$ was a periodic sequence, supposed known, with a period of $41 \times 31 = 1271$. Lay it against $\Delta(C_1 + C_2)$ in each of the 1271 possible relative positions and count the number of agreements for each one. One position should give significantly more agreements than any of the others, and that one would give the correct settings of χ_1 and χ_2. The procedure was not to be recommended as a hand

method but no doubt our electrical engineers could find a way of mechanizing it. One hoped that the method would work for other pairs so that the settings of all five χ-wheels could be found.

I went into Gerry Morgan's office to tell of these results. Max Newman was there. They began to tell me, enthusiastically, about the current state of their own investigations. When I had an opportunity to speak I said, rather brashly, "Now my method is much simpler". They demanded a description. I must say they were rapidly converted. The Research Section urged the adoption of the "Statistical Method" of wheel-setting.

Soon the electrical engineers produced the necessary machines. We called them Heath-Robinson's after a well-known cartoonist who drew ludicrous mechanical devices. The engineers were pleased because they were rushing teleprinter tape through the machinery at unprecedented speeds without ever breaking it. Well, hardly ever. Agreements, I was told, were counted photoelectrically. When you see a photograph of a contraption with teleprinter tapes running on pulleys you can be sure that it is comparing $\Delta(C_1 + C_j)$ with $\Delta(\chi_i + \chi_j)$ for some i and j.

When the five χ-wheels were all set the χ-key was added to, that is subtracted from, the cipher message to yield the combination $P + \psi$. Now we ran into the problem of "depsiing", that is separating $P + \psi$ into P and ψ. The process was analogous to depth-reading. ψ was not plain text but it had its own peculiarities. It had a great many repeated letters. Of two consecutive letters it could be said that if they were different they were likely to differ by more than the random expectation. When depsiing had gone far enough for the ψ-wheels to be set the process simplified. As a rule there would be only two alternatives for the next ψ-letter. But sometimes confusion would be caused by a repeated letter of ψ that unexpectedly covered a ψ-advance. Most of the depsiing was done by women of the Auxiliary Services, who became expert at it. When they were unsuccessful they spoke instead of "deep-sighing". When a long enough stretch of ψ was known, the motor wheels could be set and then further deciphering was routine.

The Statistical Method worked best on long messages. Fortunately the Germans were sending longer and longer messages at this time. They would put several actual messages into a single transmission. From our point of view that would be a single very, very long

message.

Production increased again. But somewhere around this time the Germans decided to change their wheel-patterns every day. Perhaps they were hearkening again to that same inspector. But they could not be relied upon to send a depth every day.

VI The coming of COLOSSUS

I come to a time now when the Research Section had given birth to two new Sections. There was the Testery, under Major Ralph Tester, where they did the actual reading of FISH messages. There was the Newmanry under Max Newman. There they were concerned with the mechanization of FISH-breaking procedures. Gerry Morgan was now Major Morgan, and John Tiltman had become Brigadier Tiltman.

I began to dream of generalizing the Statistical Method. Could it be improved so far as to find unknown wheel-patterns instead of merely setting known ones? Even if that could be done in theory it would surely need a very long message? There on my work table was a cipher message about 15,000 letters long. I must, as the current idiom went, "go to it."

First I went along the message, marking off its letters in numbered blocks of, I think, twenty. Never must the 15,000 get out of step. Never must I mistake letter number 11,614 for letter number 11,615. I then had to make the familiar statistical assumption that $\Delta(C_1 + C_2)$ agreed, sufficiently often for my purpose, with $\Delta(\chi_1 + \chi_2)$. I had to suppose that effectively $\Delta(C_1 + C_2)$ was a badly damaged version of $\Delta(\chi_1 + \chi_2)$. I hoped that it would still be at least 55% correct. Probably I could use the fact that $\Delta(\chi_1 + \chi_2)$ was the sum of two periodic sequences $\Delta\chi_1$ and $\Delta\chi_2$, of periods 41 and 31 respectively.

The first step seemed clear. I wrote $\Delta(C_1 + C_2)$ diagonally into a rectangle of 41 rows and 31 columns. (Or perhaps it was 31 rows and 41 columns? No matter). With each successive symbol the sequence would advance one row and one column. The first row was deemed to be the successor of the last, and similarly with the columns. The first symbol of $\Delta(C_1 + C_2)$ went into the first row and first column, and the sequence returned to any given square in just 1271 steps.

At the end of this procedure there were 11 or 12 entries in each little square of the rectangle, less a few corresponding to doubtful letters of the cipher text. All the entries in a particular little square corresponded to the same settings of the first and second χ-wheels. Each came in the same place of the period of $\Delta\chi_1$, and the same place of the period of $\Delta\chi_2$. Since a symbol of $\Delta(C_1 + C_2)$ was more likely than not to agree with the corresponding symbol of $\Delta(\chi_1 + \chi_2)$ I took a majority verdict in each square. Whichever symbol dot or cross occurred most often there I took to be, most likely, the symbol of $\Delta(\chi_1 + \chi_2)$ for that square. True there were disappointing cases in which the majority was only one. There were even a few in which there were just as many dots as crosses and I could make no decision. But I felt sure that my evaluation of $\Delta(\chi_1 + \chi_1)$ was much more than 55% correct. So I made a new rectangle with my estimates of $\Delta(\chi_1 + \chi_2)$ replacing the original entries.

I now had to decompose the estimated $\Delta(\chi_1 + \chi_2)$ into $\Delta\chi_1$ and $\Delta\chi_2$. To do this I chose a row in which I thought the evidence was unusually strong, and adopted it as a first approximation to $\Delta\chi_1$. Comparing this with each row in turn I could estimate whether the corresponding symbol of $\Delta\chi_2$ was dot or cross. (Dot for good agreement, cross for marked disagreement). Thus I got my first approximation of $\Delta\chi_2$. Comparing this with each column in turn I got a second approximation to $\Delta\chi_1$. Continuing in this way, alternating between rows and columns until the patterns changed no more, I got what I thought were pretty reliable estimates of $\Delta\chi_1$ and $\Delta\chi_2$. True there were one or two entries in these for which the evidence was very weak. Moreover there was the usual ambiguity; did I have the correct patterns or their reverses? For deltaed sequences this matters.

There were two considerations that allowed me to do some fine tuning. $\Delta\chi_1$ and $\Delta\chi_2$, merely because they were deltaed periodic sequences, would each have to have an even number of crosses. Their reverses, those imposters, would each have an odd number. Moreover the χ_i-sequence had to be plausible. So I got estimates of $\Delta\chi_1$ and $\Delta\chi_2$, and thence of χ_1 and χ_2, that I thought reliable. At least I was content to assume their accuracy as a working hypothesis. In the event I did not have to correct them.

I now had to go through the same process again with a different pair of impulses.

It was to be feared that the agreement between $\Delta(C_i + C_j)$ and $\Delta(\chi_i + \chi_j)$ would be smaller. On the other hand by making $\Delta\chi_1$ or $\Delta\chi_2$ one member of the pair we could take one of the patterns $\Delta\chi_i$ and $\Delta\chi_j$ as known. In fact I got satisfactory results with my new pairs and soon had the five χ-patterns. Barring any residual errors I now had the χ-key. Accordingly I "dechied" the message, getting $P + \psi$. Now for a grand operation of depsiing!

In my deciphering I was at a disadvantage compared with the wheel setters, since I did not know the patterns of the ψ-wheels. But I made some progress and it soon became clear to me that I had made a lucky choice of message. The text seemed to be in short blocks separated by unnecessarily long sequences of symbols of punctuation. At any rate the sender had been generous in providing blocks of 4 or more word-spaces. I wish I could remember more about how I separated P from ψ. I know I made an effort to construct the motor-key from the ψ-motion as I went along. In this year 2000 I see clearly that ψ-motion would repeat on the period of the motor key, that is $61 \times 37 = 2257$, and this fact should have been exploitable. If there was here a ψ-letter repeated twice there would likely be the same 2257 places on. If so one would have 32 possibilities for a triplet of plain text there, and only one might be plausible. I was equally capable of seeing this when I was working on that message. But whether I used this observation or neglected it in favour of some procedure that I have now forgotten I cannot say. All I can be sure of now is that I did get enough information about ψ, its stoppages and its starts, to determine the patterns of the ψ-wheels and the motor wheels. What was left of the message could then be routinely deciphered.

I remembered taking this message, with its wheel-patterns and wheel-settings, to the Testery. To emphasize what I considered the importance of the occasion I told Ralph Tester "This is the first machine to be broken on a depth of one."

Statisticians improved upon this "Rectangle Method". Instead of taking the majority verdict for a little square they estimated the probability of $\Delta(\chi_1 + \chi_2)$ being cross, and wrote that number there. In getting the first approximation to $\Delta\chi_2$ they would not make a crude estimate of dot or cross but work out probabilities. And so on for further approximations. The process became that of finding an eigenvector for a rectangle-matrix.

More wheel-patterns were found.

Clearly the refined Rectangle Method would be too tedious to be done routinely by hand. If it was to be of real use mechanization would be necessary. However our Post Office engineers were steadily improving their FISH-breaking apparatus and they went on to design an electrical machine that could find χ-patterns by the Rectangle Method. That was COLOSSUS, the machine that is now recognized as an electronic computer, and which has been claimed by Bletchleyites to be the first of that race.

I should remark here that in those days the telephone system belonged to the Post Office and so the engineers who looked after it were Post Office employees. It was they who built COLOSSUS from off-the-shelf units.

I remember being introduced to COLOSSUS. With other members of the Research Section, I was taken to a large room, where a large box-shaped object, sheathed in sheet metal, stood upon a wet floor. If it was $16 \times 3 \times 5$ in feet, that would not contradict my memory. "That" we were told "is COLOSSUS. Its inner circuitry contains 1500 valves." (a.k.a. vacuum tubes). Gerry Morgan, gazing at the wet floor, remarked that it had not been house-trained yet. We were told that those valves generated heat and the apparatus had to be water-cooled. Alas, there was some leakage.

That memory now puzzles me. It does not fit the COLOSSUS now installed at Bletchley. There the valves are arranged on racks, exposed to the open air and the public gaze. And there are whirling teleprinter tapes to draw the visitors' attention. That suggests wheel-setting, whereas the chief glory of COLOSSUS was wheel-breaking. I acknowledge that most of the differences could be explained by the engineers deciding that air-cooling was better than water-cooling. I regret that I was too theoretical a mathematician to study and remember the appearance of the Newmanry's machinery.

One of my memories of early COLOSSUS-time is still vivid. It is of Max Newman exclaiming with an air of surprise "You know, this thing could do logical operations!" (Silent upon a peak in Darien).

I cannot remember any further work I did with FISH. I knew of various complications the Germans now put into their cipher machines. They took the form of additions to the motor key. Sometimes a χ- or ψ-impulse from a few letters back would be added to it.

Sometimes even a delayed impulse of plain text would be added. The latter device was called an autoclave. From the German point of view it had the advantage of making depth-reading impossible. On the other hand a single wrong letter could wreck the remainder of the message. It seemed that the autoclave, when used, gave more trouble to the Germans than to the Bletchleyites. These new devices did not prevent the breaking of the χ-wheels, but they made depsiing and ψ-breaking more difficult. On the whole however Bletchley kept on top of FISH until the end of the German War, armed as it was with COLOSSUS. Not all messages were read, but there was a statistical test for deciding early on whether the one under investigation was likely or not to yield. If not, then it was rejected.

VII Comments

Sometimes in phantasy I envisage a future University Department in which, for reasons unfathomable to this Twentieth Century,[2] they are devoted to the study of how FISH was broken and, more importantly, of how it ought to have been broken. There is Professor Cumlatly, lecturing at his Annual Cryptographical Meeting of 2253. I cannot distinguish what he is saying, though a few phrases sound like "Turing's Method", "Rectangle Method" and "HQIBPEXEZMUG-Method". He seems to be pouring scorn on all of them. One sentence comes through quite clearly. "and yet STURGEON could have been broken so easily." There is another brief interval of clarity at Question Time. Someone asks "Johnny, wouldn't it have been quicker and easier to break those keys from the depths and near depths of 1942 by the Rectangle Method?" The Professor answers "That is a good question. I already have a graduate student working on it."

In this phantasy I have expressed my fear that cryptographic historians, having mulled over their problems for a few years, will express pained surprise at the cryptoanalytic crudities of the Nineteen Forties and will marvel at the obtuseness of their practitioners. If so let them remember that our customers wanted results quickly; the first method that gave any results at all would have to go right away into general use.

I put in a reference to STURGEON, the teleprinter cipher machine that I have said

[2]It takes 2000 full years to complete 20 centuries.

we did not break. Like TUNNY it had two subkeys but now they were generated by ten regularly moving wheels. The first, call it the χ-key, was added to the plain text to produce a half-encipherment $P + \chi$. The second subkey permuted the impulses of the half-enciphered letter. There were variations but typically the first and second impulses would be interchanged if the second subkey had a cross in the first impulse, but would be left unchanged if that key had a dot there. The next impulse of the subkey would act similarly on the second and third impulses, and so on.

STURGEON's physical machine, it is now known, had one disastrous weakness. There was a switch that put the wheels back to their original setting. When an operator had sent a message he was tempted to use that switch and so send the next message on the same key. It saved him the trouble of turning the wheels to a new setting. So depths of six or seven, or even more became common in the STURGEON traffic. With depths of that profundity messages could be read, permutation or no permutation, wheel-patterns could be deduced and the working of the machine laid bare. Deep depths were common enough to yield a supply of decodes. So it could be said that STURGEON was broken. But to workers on TUNNY and its analogues a month would not be considered broken until we had some way of decoding messages that were not in depth, as by dragging, Wyllie's Method, the Statistical Method of wheel-setting or the Rectangle Method. We never got that far with STURGEON. Eventually it was decided at a high level that the flow of information from STURGEON was less satisfactory, in volume and in content, than that from TUNNY-type traffic. We were instructed to concentrate on the latter. In spite of its fantastic misuse by its German operators I can think of STURGEON as "the one that got away". In my phantasy I express my misgivings that perhaps, by some method that we missed, STURGEON could have been made as tame as TUNNY.

The question and answer that I put into my phantasy record a question I have often asked myself. But let the answer be Cumlatly's own. Having got a key K we can write

$$\Delta(K_1 + K_2) = \Delta(\chi_1 + \chi_2) + \Delta(\psi_1 + \psi_2)$$

But $\Delta(\psi_1 + \psi_2)$ is about 70% dot. Hence $\Delta(K_1 + K_2)$ must be in 70% agreement with $\Delta(\chi_1 + \chi_2)$. And the agreement would still be 70% for other pairs of impulses. Surely with 1000 or so letters of key the Rectangle Method would have resolved $\Delta(K_1 + K_2)$

into $\Delta(\chi_1)$ and $\Delta(\chi_2)$? I suppose I could check this if I went to enough trouble. But what would be the point now, so long after the last FISH message was sent? On second thoughts the problem could be reformulated as one in Pure Mathematics and so made worth solving. It would not be the first of Bletchley's little problems to undergo such a transformation.

If I were asked to state the main weaknesses of the FISH machine I would comment first on the separation of the five impulses, each being effectively controlled by only two wheels. Actually the two motor wheels made four but our cryptoanalytic methods did not have to consider the motor patterns in their early stages. In STURGEON, because of the permutations, all ten wheels were involved in the production of a single impulse of cipher text. The second main weakness was that the ψ-wheels moved in step. Had they moved independently, each with its own motor-pattern, then Turing's Method would not have worked, and the Statistical Method and the Rectangle Method would not have applied.