# :: Snerx's N-Dimensional Cipher ::

**ABSTRACT:** Sniph is a dynamically shifting multi-dimensional cipher. The impulse for creating this cipher was twofold: to create the first scalable N-dimensional cipher and to create the first cipher that is *in principle* impossible to knowingly bruteforce. This is accomplished by using a combination of several traditional cryptographic techniques as well as a new one termed library surfing, with segmentation between the two that makes things like perfect forward secrecy obsolete since the keys can never be reverse-engineered without creating a hermeneutical lacuna, or gap in the knower's capacity to know which keys were used with the original plaintext. Further, I wanted to be able to use a homebrewed cipher without higher-maths so the entire process could be easily done by hand.

## :: Introduction and Meta-Epistemic Concerns ::

Sniph started as a project to create a sufficiently-difficult-to-crack cipher that could be done by hand on paper for the use of sending encoded messages between friends in board games. This means that this is all done without the need for higher maths or any computational machines. Big claims from a dilettante, I know, but hopefully the underlined current public code and following descriptions substantiate these Icarusian claims.

I am not a cryptographer, and while cryptography is fascinating to me, I am not good at it. I gave and received a lot of opprobrium about my cipher during its initial conceptual incubation on lainchan.org because I was claiming to have figured out something substantial that no one who knows better has figured out. In defense, I am a trained logician and have considerable background in meta-epistemology; this is how one can make the claim that one is able to determine some advanced properties of cryptographic systems without knowing anything about them prior. But I am still stupid and have probably missed something obvious and important and devastating to the whole project. Because of this, I am open to edits and revisions of this document should the cipher be shown to be insecure or bastardization of terminology is found.

I am anticipating someone will find a way to break this cipher, but I stand by my statements to come later: why bother doing cryptography anymore unless you think there is substantial room to improve? Why bother doing anything unless you want to be the best at it, or at least solve a fundamental problem that has plagued everyone else up until that point? I hope that in discovering problems with Sniph, potential solutions to those problems will also be discovered, and thereby solutions to perennial epistemic problems will be evinced.

I want to open by limning exactly what is meant by 'in principle' when I say Sniph is impossible to knowingly bruteforce *in principle*. Epistemically speaking, nothing is locked away 'in principle'. What is meant by this is that there is no possible information that is locked away from any knower for principled reasons; only for practical reasons could information be locked away, further that this is *necessarily* the case. Wittgenstein formally proved this with his private language arguments which showcase that any principle gap or divide in information between knowers results in an ultimate and total divide between

the knower and all possible knowledge.[1] This is to say private language leads to a complete collapse of any epistemic framework you want to posit - it makes knowledge itself impossible.

Many people have a hard time understanding Wittgenstein's private language arguments, and my two-sentence summary may be disagreeable to some, so instead I will give a far simpler original epistemic example that illustrates this point. I like to call this the *opaque box* example.

Let's say there is a box somewhere in the universe that is totally opaque, not just to light, but to any possible probing or analysis. We cannot know what is inside this box, and not just for practical reasons like lack of proper tools (similar to what José Medina termed a *hermeneutical lacuna*),[2] but for principled reasons like lack of *any* possible understanding. Okay, so we can't know what's inside the box, and we can't know in principle, it's locked away from all possible knowers. Anything then can be inside this box, and anything means *anything*. This includes information that might make knowledge itself impossible. The box could contain information or conditions that make its own existence, or even potency for existence, impossible. This means the opaque box self-defeats. It is an impossible object, like a square-circle; the thing allows the impossibility of itself and also forces its own denial, therefore you do not have that thing, and necessarily so. This is apodictic, so anyone who argues to the counter can be dismissed as misunderstanding what it is they argue for.

Okay, fine you say, but what does this have to do with cryptography? Well we know that information can only be locked away for practical reasons, which means we now know that *all* ciphers and encryption algorithms work off of *when* not *if*. This is to say that all ciphers must necessarily be concerned with the amount of time it takes to bruteforce them, not if they can be bruteforced in the first place. Some might point out OTPs as an exception to this, that if each pad is truly one-time, then it would be impossible to decipher the ciphertext because the keys used to establish it are truly unique, or *truly random*. Claude Shannon was the first person to mathematically prove this impossibility, but it should be noted that math is second-order logic, and is subservient to first-order logic. In the same work Shannon illustrates many other important properties of cryptographic systems like problems with symmetric keys concerning the use of the same key to cipher as well as decipher, and so on.[3] This is all notable for later, as maths operate a narrow domain when compared to formal logic, and these problems of symmetric systems are only problems if symmetry is maintained in this narrow sense.

If OTPs can do what they have just been described as doing, they should be wildly useful because they would allow for a *principled* lock-away of information, where no amount of probing could ever recover the plaintext. But everyone reading this now knows that principled lock-aways are not possible given the meta-epistemic conditions for knowledge. So OTPs are necessarily stinted in their use, and this partly explains why they haven't caught on: if it's truly random, then you need to somehow communicate a new keyset for every new ciphertext, which becomes very impractical very fast and requires automation, stealing control of the keys away from human use (problematic because we still want to be able to do this on paper).

As a potential solution to this problem, there are ciphers that implement key streaming, which is meant to allow for wrapped or modular addition of characters, as if to make an OTP infinite in its list of keys. This treats the symptom but not the problem itself, because if you know the PRNG source that is used to deterministically generate those keys, then you can break the OTP just the same since the 'password' is effectively the PRNG source. 'True' randomness is possible if you want to use robust CSPRNGs (like Cloudflare's LavaRand[4]), but this is highly impractical for one-off messages, or the freezing of information for documents like encrypted PDF's, or if you don't have an internet connection, or if you

[1] Specifically from remarks 243-304 in *Philosophical Investigations* - https://snerx.com/archive/Wittgenstein-PhilosophicalInvestigations.pdf
[2] José Medina, *The Epistemology of Resistance*, Oxford: Oxford University Press, 2012.
[3] Shannon, Claude. "Communication Theory of Secrecy Systems". *Bell System Technical Journal*, 1949. 28 (4): 656–715. doi:10.1002/j.1538-7305.1949.tb00928.x
[4] This is a good read if you weren't previously aware it existed - https://blog.cloudflare.com/randomness-101-lavarand-in-production/

want to have direct control and access to the keys being used to encrypt the data (again, automation ruins all of this).

An additional potential solution is asymmetrical cryptography (public-key cryptography), which does solve some issues, but not the fundamental issue of bruteforcing. Once a private key is bruteforced, all other messages signed with the same key are now completely open. However, asymmetrical cryptography is so computation-intensive that it is *infeasible* to bruteforce even if given millions of years.

This is not good enough for me, and it shouldn't be for you either, not just because there may be sufficient computational devices invented in the future that are orders of magnitude more powerful, like quantum computers (I have seen some mixed debates as to whether these are actually devastating to current encryption algorithms), but because we are looking to cross the divide between the practical and the principle here. Millions of years to *exhaust* a list of keys is not millions of years to find one randomly - people get lucky, PGP has been cheated before, as well as other private key systems[5], and it will happen again.

So the goal is that even with infinite computing power and all possible keys given instantly, you could still not tell if you've cracked the cipher. And this should always be the goal from here on out in cryptography; why bother doing cryptography anymore unless you think there is substantial room to improve in this regard? For that matter, why bother doing anything unless you want to do it better than anyone else ever will? Touch the sun **and** keep your wings.

Sniph is an attempt at collapsing the distinction between principle and practical modalities at all levels, not just in mathematics for cryptographic purposes, but for meta-epistemic thought experiments as well. If this is successful in the way I describe, Sniph will serve as an incredibly useful example for meta-epistemic discussion surrounding modalities and general properties of knowers in epistemic systems on top of its use in cryptosystems.

I try to accomplish this by doubling down on the problems with OTPs and symmetric key systems in general; I believe that when stacked a certain way their problems can negate each other (features, not bugs). This leads to what I term 'library surfing' and is based on a similar conceptual structure to that of the search algorithm for the Library of Babel's website.[6] I've been told that important parts of this may be considered similar, but not the same as, raw XOR cipher functionality, but a significant difference between Sniph and a XOR cipher is that in Sniph you cannot add/subtract/XOR the plaintext, ciphertext, or keys from one to reveal the other. All these properties combined make Sniph immune to plaintext attacks. All this combined also gives Sniph the property of perfect secrecy without the need of communicating a new keyset for every message.

These seemingly fantastical statements may be off-putting to cryptographers, for, "Few false ideas have more firmly gripped the minds of so many intelligent people than the one that, if they just tried, they could invent a cipher that no one could break."[7] But this is not a matter of merely gaining a 'sufficiently large keyspace', but of gaining an *infinite* keyspace, a topologically recursive keyspace, an indefinitely traversable keyspace, with no way of knowing if any of the space traversed is meaningful space unless you were the one who assigned the keys in the first place.

So when I say Sniph is impossible to bruteforce *in principle*, I am saying that it is impossible to determine *meaningful keys* qua bruteforce attack. I understand this may be an unfair use of the term 'in principle' here since the relevant keys can in fact be bruteforced, and are expected to be bruteforced, but this technical abduction of the term allows us to speak on the interpretive gap the bruteforcing process creates without having to rely on purely meta-epistemic terminology like the earlier referenced *hermeneutical lacuna*.

[5] For funsies, here's a list of Bitcoin addresses with money on them that had their private keys bruteforced - https://lbc.cryptoguru.org/about
[6] The principle the search function operates on is what I believe the most direct parallels to Sniph can be drawn to - https://libraryofbabel.info/
[7] I am quoting David Kahn, from his book *The Codebreakers*.

An outline of how the protective property of Sniph works: If you obtain a plaintext and paired ciphertext, you will be able to traverse all possible keys within the given space defined by the length of the ciphertext (the length of the ciphertext effectively determines the analogous page-length in the Library of Babel, this is explained in detail in the next section). Every single one of those keys will be able to produce *all possible* plaintexts from *all possible* additional ciphertexts, because in traversing this keyspace you will have recursively iterated *all possible keys*. This means having plaintexts with paired ciphertexts will not reliably produce the keys you desire. Unintuitively, a more accurate way of saying this same thing - it will *always* produce the correct keys, but they will be mixed in with all other possible keys that all work equally well and so you will never be able to know which keys were the original ones you desired, only that they were definitely produced.

As a final note to the introduction, this is not just some academic project. If this cipher, or any cipher that fulfills the epistemic properties I have already described work as I say (that is, crosses this principle-practical divide), then we will have the first case in history of an 'epistemic blackbox'. This is the only way, in principle, to guarantee the basic human right of privacy; that without a principle lock-away there is no means by which to ground any claims of necessity when it comes to saying no one except those intended can access certain information. This then attains ethical imperative. This project, or one dangerously similar to it, must be finished else we necessarily perverse our capacities for saying anything can be private - as without this, nothing is actually private. As such, if you find a system-breaking flaw with this cipher, you ought to make it known, lest you perverse your own rights.

If you agree, or even if you hated reading this whole thing and think I'm terrible, feel free to yell at me here - snax@snerx.com.

**:: Process and Properties of the Cipher ::**

**OVERVIEW:** The way Sniph works is by generating tables of variable sizes (like Polybius squares) that are then seated inside of each other to the N*th* dimension specified either by the user's passphrase/key or manually by the user. The breadcrumb/pathway through each dimension/table-set is then determined by the passphrase. This process, with additional minor steps, is repeated for every single letter of plaintext fed through the cipher giving it similar, but not same, properties of one-time pads. Seating tables of variable sizes gives this cipher the additional property of having one-to-many inputs-to-outputs and one-to-many outputs-to-inputs, but only sometimes many-to-many inputs-to-outputs. The ciphertext generated and this property of it can be thought of as being similar to a dynamic version of the Spanish Strip Cipher.[8] This directional one-to-many property helps substantially with obfuscating data; when this property is stacked, this is what I term *library surfing*.

**IN PRACTICE:** You start by selecting an arbitrary plaintext to cipher, like 'HELLO', or 'WORLD', and by selecting an arbitrary passphrase like, 'ALICE', or 'BOB'. After this, the first thing that the cipher needs to do is generate the tables; the cipher decides what table dimensions to use either by flags sent in the CLI, manual entry in the CLI, or by the given passphrase. For this example we'll use the passphrase option so we can see how the process works in its intended mode.

       Given the passphrase 'BOB', the cipher deterministically generates the table width, height, and then depth. To do this, we look at the first letter of the passphrase 'B' and take the numerical relation of that letter '2' and use this to decide that the table width will be 2. Then we look at the second letter 'O' and take this to mean the table height will be 15. Following the same procedure, the depth is then 2. Due to the restraints placed on table dimensions by the charset being used (understood later), the minimum table dimensions that can be used are 4x4x4. This means that we will instead be adding each of the numbers we just found to 4. The table dimensions are then 6x19x6. While the cipher itself is N-dimensional, we still want to be able to easily do this by hand on paper, so we will be wrapping any numbers over 10 so that our tables on our graph paper don't become too large. The table dimensions are finally 6x9x6.

       Because N-dimensions means *N-dimensions*, you could manually set the table dimensions to be *1,000 x 10,000 x 100,000*, but we wouldn't be able to do that by hand, and so anything over 10x10xN is currently not implemented in the code (but will be).

       Now we have to populate the tables with characters so that a plaintext character can be chosen. Since these tables have depth, or *seated dimensionality*, and because drawing and filling out a cube on paper is difficult, we will use the table depth (all the stacked 2D layers) up until the last two layers (*N* minus 2) merely as a means to produce a Caesar shift in our cipher.

       To do this, we look at the first table, 6x9x**1**, then 6x9x**2**, then 6x9x**3**, then stop at 6x9x**4**. The first table is 6x9, and to determine which cell to open, we again turn to the passphrase. In this particular case, all the character spots of the passphrase have already been used, so we will wrap the passphrase over to give us additional characters (streaming the passphrase into 'BOBBOB'). The passphrase's fourth letter's numerical relation is 2, so we go 2 across from the left of our initial table, giving us the second column. The passphrase's fifth letter's numerical relation is 15, so we go 15 down from the top of the second column, giving us the sixth row. We represent this cell as '2-6'. We repeat this process for the second 6x9 table opened, using the sixth and seventh characters from the passphrase 'BOBBO**BB**OB', giving us '2-2'. The same with the third table, '6-2'. And with a final passphrase of 'BOBBOBBOB**BO**B', the

---
[8] Mystery Twister has a good summary and implementation - https://www.mysterytwisterc3.org/images/challenges/mtc3-sanguino-03-SSC-03-en.pdf

fourth table, '2-6'. These are the *path-tables* and are used to determine the Caesar shift for the final two *end-tables*.

Now that we've determined the path-tables, we will sum each table's picked cell and then add each sum together for a total which will be used to shift the character sets in the end-tables. We start with the first table of '2-6', which gives us -4. The second gives 0, the third 4, and the fourth -4. Adding these all together we get negative four ( -4 + 0 + 4 + -4 = -4 ). You may notice that working with a very short passphrase and lots of repeating numbers gets us nice small results to work with all the way through, but of course this is arbitrary in the code and passphrases with higher entropy result in more dynamic numbers throughout this process.

We can populate the end-tables now. This is visually represented in the original gdoc I made about Sniph. The second-to-last end-table's first cell gives us our first last end-table. I am aware of how poorly named these things are, just bare with me. To reiterate, in this example there is a fifth 6x9 table in the total of six sets of 6x9 tables; this table is our first 'end-table', or our second-to-last table total. Picking any cell in that table opens up another 6x9 table, which is our second 'end-table', or our last table total. Hopefully this clarifies it. The gdoc is there to help if imagining this gets too tricky.

So we can look at the first cell in our second-to-last table here, which would open a final 6x9 table. This first-cell's final-table's first cell would be populated by the first character in our character set, and for this example our character set will be 'ABCDEFGHIJKLMNOPQRSTUVWXYZ'. That makes the first cell A. Going across the row to the second cell, we would see B, then C, then D, then E, and finally F, which ends the row. The second row starts with G, and the character set continues to wrap row by row until it runs out. Again, this is visualized in the linked gdoc. In a 6x9 table there are 54 cells, which means at the 27*th* cell, where the character set would otherwise be finished, we wrap the character set again and start at A, then B, and so on. At the 53*rd* cell we would begin wrapping the character set yet again. This means this table's last cell (the 54*th* cell) ends with the character 'B'.

Opening the second cell in our second-to-last table here, we would find another final 6x9 table. This second-cell's final-table's first cell would be populated by the character now wrapped over from the end of the first final table, which in this case is 'C'. The second cell in this final table would be D, and so on. The character set keeps wrapping and continuing over and over again until each final 6x9 table is populated. At least, this is how it would be if we started the first cell of the first final table with the first character of the charset.

Instead, our path-tables told us that the character set was to be shifted back four characters, and so instead our charset looks like 'WXYZABCDEFGHIJKLMNOPQRSTUV'. This means the table-populating process is carried out the same way, but with different starting characters. This is how the Caesar shift is applied in Sniph.

Why split up the end-tables into two sets of equally sized tables? Why wrap the charset over and over if you only need one character to be chosen? This property of Sniph, using seated tables, allows us to arbitrarily pick one of many possible outputs to code for our plaintext character. This is homophonic substitution and means that encoding the same plaintext character over and over using the exact same passphrase, the same pathways and table dimensions, will generate many different ciphertexts. This makes the re-use of a symmetric key far more secure since there are no consistent outputs that can be amalgamated.

Back to this in action, if our plaintext that we want to cipher is 'WORLD', and our shifted character set has populated the first final table, we now have three choices of cells to pick that contain the character 'W'. There are two more choices in the second final table, and so on. We can pick any of these cells to represent 'W', the choice is arbitrary. In the code, this is chosen at random. On paper, the choice is at the pen's discretion.

Let's say we picked the very first cell of the first-final-table, this would be cell '1-1' of the second-to-last end-table, and cell '1-1' of the first-final-table. This is then the ciphertext '1111' that codes for the

plaintext 'W'. If we had picked a W from the second final table '1-2', at cell '5-1' for example, then the ciphertext would be '1251'. There are of course many other cell locations that code for 'W' in this example. *N.b.*, having many outputs for a single input is not sufficient for adding the kind of difficultly required to make this cipher truly secure. Rather, the final transformation of the output as ciphertext is what makes this property notable.

      If the end-table's coordinates were always direct like this, where the first cell of the first table is always '1111', this would make it obvious as to which cell of which table was selected regardless of the table dimensions. This also means that if a long ciphertext never uses a number above 4, then an adversary intercepting the message would know the table dimensions were 4x4. To stop this kind of analysis, the un-used numbers for a table dimension are wrapped as stand-ins for the already used numbers. In a table-size of 6x9, since the rows are only 6 long, 7 can stand in for 1, 8 for 2, 9 for 3, and 0 for 4. Since the columns are only 9 long, 0 can stand in for 1 there. *E.g.*, our first choice for 'W' as '1111' could also be represented as '7070' or '7110' or some other combination. This is also chosen at random.

      To make this clear, not only are the cell choices many/arbitrary but the cell *representations* are also many/arbitrary, at least for table-sizes less than 10x10. This stops analysis of ciphertext from being able to determine the table sizes being used, since all numbers for all sizes are always present in all ciphertexts. Additionally, the ciphertext only ever tells you the cell locations being chosen for the end-tables, they don't include any information regarding table depth, the seated dimensionality, so the Caesar shift cannot be determined from analysis of the ciphertext.

      This is uncommon for a cipher, since most ciphertexts that have been Caesar-shifted can be directly shifted without damaging the process for returning the plaintext message, whereas this is not the case with Sniph. Directly shifting the ciphertext output from Sniph ruins the possibility of being able to reverse the ciphertext into plaintext since the ciphertext output are coordinates whose locations only make sense qua the original Caesar shift. An even bigger claim - this property of Sniph makes it possible to have both random **and** predetermined keys in the same cipher. I will say this again but in a different way - this allows you to use 'unique OTP keys' without having to transmit the new keys each time you generate ciphertext; this allows you to reuse the same passphrase and not generate similar-looking output each time. This solves the problem with the usage of one-time pads not being one-time.

      What we end up with here is four numbers coding for a single character. We have our four-character long ciphertext for our single-character plaintext. So all plaintext ran through Sniph is expanded by a factor of four. For almost a year this remained the case, but we now replace those four digits with their respective character encoding given the UTF-8 standard ('1111' returns as 'ω')[9], and so Sniph now has 1:1 input-to-output ratio which keeps data size down.

      The more perspicacious amongst you may have realized that the same principles of Sniph can work for a single table with a Caesar shift applied, and so we would only need to expand text by a factor of two, however I made the original pen-on-paper cipher when I was 13 and didn't think this far ahead. In the future, the code will probably reflect more efficient uses of the cipher in this way, but for now it works the way it works.

      That's our first character of plaintext. If we moved to the second character of plaintext now, we would repeat the entire process described up until this point. When doing this by hand on paper, I have been re-using the same table dimensions for all plaintext characters to save time, however in the code for Sniph this is not the case, instead it re-determines table sizes per-character of plaintext. This means that while our first character got tables of the 6x9 size, our second character would get something different (or possibly the same again if the passphrase happened to wrap around starting over on 'BOB').

---

[9] A complete list of characters in UTF-8 with the 1,111th character being displayed first - https://www.fileformat.info/info/charset/UTF-8/list.htm?start=1111

This repeats for each character being ciphered, and the passphrase is streamed, or continuously wrapped, until all characters have finished being encoded.

**BREAKING SNIPH:** To understand this in more detail, and to see why this is so devastating to attackers trying to bruteforce Sniph, we can quickly try to reverse this process.

Let's say we were an attacker, and we knew how Sniph worked, and we intercepted the ciphertext '16671251565263370190' (v0.5 of Sniph, which is not secure, output this from the plaintext of 'WORLD'). Okay, we want to be able to determine the plaintext that the ciphertext codes for, so we need to guess the passphrase used to determine the table dimensions. We bruteforce the ciphertext starting with 'AAA' and crawl through every possible permutation, ending with 'ZZZ'. Given current standard commercial processing power, this should only take a few seconds to bruteforce. The bruteforce program immediately shows the entire list of characters that '16671251565263370190' could code for given each possible passphrase used. One of the passphrases bruteforced was 'BOB', which was indeed the original passphrase used (given that it wraps). However, every possible five-letter plaintext was generated in this list given every possible three-letter passphrase. This becomes very problematic very fast.

Something interesting happens here. Yes, you will have generated the correct plaintext after hitting the correct passphrase of 'BOB', but you will have generated the correct plaintext with many other passphrases in the permutations as well. Many passphrases will get you the desired plaintext, however only one will get you the desired plaintext in the desired table dimensions. So there are two things going on here as this is the second time the one-to-many property is functioning in Sniph. The first was in selecting the coordinates for the ciphertext, the second here is in selecting passphrases to match ciphertexts. Since the table dimensions do not have to be determined by the passphrases, the CLI user or paper user can arbitrarily select their own table dimensions, making passphrase bruteforcing an order of magnitude more difficult (I believe it actually becomes six orders of magnitude more difficult, will need to check with someone better at math).

So this time, in order to find the passphrase used with an independently chosen table dimensionality, we will have to try every passphrase permutation (from 'AAA' to 'ZZZ') for every possible dimensionality. But that's fine we say, because we can still spit out all the possible permutations in a couple minutes. So we run the bruteforce program again and get a list of plaintexts generated by passphrase permutations 'AAA' to 'ZZZ' for the 4x4x4 table dimension. We get another list for the 4x4x5 table dimension, and another for the 4x4x6 table dimension. Wait, how many seated dimensions do we try?

Well, with a 26-character-long charset, the total number of Caesar shifts that would be meaningful is 26 shifts, since the charset becomes mirrored for each table permutation after 26 shifts. So we say we'll go out to NxNx26 tables, since meaningful shifts are not likely to happen beyond 26 transformations. Okay so we start over with 4x4x4 to 4x4x26 and then repeat the process with 4x5x4 to 4x5x26 again, all the way to 10x10x26 since we're working with base-10 numbers in the output. I'm sure someone smarter than me will be able to say how many total sets this amounts to.

Let's pretend this would still be near-instantaneous and we get our list of plaintexts paired with each possible passphrase in each dimension-set for our single ciphertext. The problem now is compounded, for by allowing *every possible dimensionality* to be explored, with *every possible passphrase*, we will have generated *every possible plaintext* in a five-character space. This is the third one-to-many property of Sniph.

This third one-to-many property of Sniph works on every single pairing of ciphertext and passphrase; you generate every possible character-combination as plaintext. This means that in our example of a twenty-character ciphertext, we will have generated every five-character latin-script text

that has ever existed, will ever exist, and can ever exist. Every possible permutation within a five-character space will have been covered, and further, it will have been covered many times over. This means that not only will we have many valid English plaintexts that can stand in for our ciphertext, but many valid N-language plaintexts that can stand in for our ciphertext, all possible valid plaintexts actually, and multiple sets of them for *each passphrase*. This means there is a recursive multiplicity to our generated sets.

For a last attempt at breaking Sniph, I think we get the most interesting results. This time we're not blindly bruteforcing a ciphertext from Sniph, instead we intercepted the original plaintext and are trying to determine the passphrase used so we can decipher future transmissions. As the attacker we perform this plaintext attack; we know the plaintext was 'WORLD', and we bruteforced the ciphertext just the same as in the prior attempt. This time we only care about passphrases that result in 'WORLD' given the ciphertext, and this leads to the fourth and final one-to-many property of Sniph.

All passphrases will always result in 'WORLD' given this ciphertext or any other ciphertext. Since the dimensionality of the tables can be determined independently of the passphrase, there will always be a possible dimensionality wherein any given passphrase can generate any plaintext from any ciphertext. You can find a valid combination of plaintext, passphrase, dimensions, and ciphertext, for any given plaintext, passphrase, dimensions, or ciphertext. This is what it means to surf the library.


To recap: You start with some arbitrary plaintext, like 'HELLO' or 'WORLD', which gets you some arbitrary ciphertext '11112222333344445555', and you try every password permutation from 'AAA' to 'ZZZ', which gets you every plaintext permutation from 'AAAAA' to 'ZZZZZ' in every dimensionality. This means a robust plaintext attack will result in many (read: all) passwords in that five-character space returning as valid passwords.

I want to briefly expand on this particular aspect here; that five-character space is what I am calling a *library*. This is because for each set character-space, there is every possible character permutation that can be crawled through, very similar to the property that the Library of Babel has. Although, the Library of Babel is only one library, its algorithm only crawls through a 2,500-character space. Sniph has infinitely many possible libraries, since any length can be used as a plaintext input for ciphering (or a ciphertext input for deciphering). So note that computational difficulty for bruteforcing Sniph is not based on how long the passphrase is, but is exponentially more difficult the longer the ciphertext is.

Library *surfing* is not used to reference the process of ciphering plaintext in Sniph, but rather to the process of deciphering ciphertext. This is because you are surfing between libraries, between possible sets of permutations, when you attempt to decipher in Sniph. This is also the source of the talk much earlier about 'hermeneutical lacunas' since there is a gap created when attempting to interpret the output from a bruteforce attack.

So one-off plaintext attacks in Sniph do not produce reliable information. Further, all passphrases that fail in one dimensionality work in some other dimensionalities, making all passphrases always work ultimately. That said, there is a potential way to break Sniph given *sustained* plaintext attacks, and this is covered in the next section.

It has been suggested that if the same passphrase is used to generate ciphertext from the same plaintext over and over, a modified frequency analysis might be able to determine the characters of the plaintext from the ciphertext, as demonstrated on the Spanish Strip Cipher by Luis Alberto Benthin Sanguino here.

The good news is that a solution to frequency analysis is already implemented in Sniph since each character gets its own set of unique homophones. But like I said much earlier, why bother doing this unless you aim to be the greatest ever, so I also made a modified version of the Spanish Strip Cipher; the number of digraphs/substitutions is proportionate to the frequency of occurrence of the character it is meant to replace. This is shown below and there is now a code implementation by OdiliTime here.

**SNERX STRIP CIPHER (SSC)** — 2018 / 09 / 15
49 x 20 = 980 — 49 x 5 = 245 — with DFV = 394

Frequency analysis dipped by higher output count for high frequency characters.

Key for table can be deterministic or table can be randomly generated and saved as the key.

Frequency variable default shown. Can be manually determined or deterministic based on key/pass.

A 26x26 Vigenère Cipher is secure if the total length of ciphertext is less than 27.67 times the length of the keyword. This 49x20 Strip Cipher is secure if the total length of ciphertext is less than ~394 times the length of the keyword without the freq var. With freq var attunement this approaches infinity; all possible plaintexts become equally valid.

| Char | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X | Y | Z | { | } | \| | : | " | < | > | ? | ~ | ! | @ | # | $ | % | ^ | & | * | ( | ) | _ | + | /n | TAB |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Set | a | b | c | d | e | f | g | h | i | j | k | l | m | n | o | p | q | r | s | t | u | v | w | x | y | z | [ | ] | \ | ; | ' | , | . | / | ` | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | - | = | | § |
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 | 32 | 33 | 34 | 35 | 36 | 37 | 38 | 39 | 40 | 41 | 42 | 43 | 44 | 45 | 46 | 47 | 48 | 49 |
| 1 | 001 | 002 | 003 | 004 | 005 | 006 | 007 | 008 | 009 | 010 | 011 | 012 | 013 | 014 | 015 | 016 | 017 | 018 | 019 | 020 | 021 | 022 | 023 | 024 | 025 | 026 | 027 | 028 | 029 | 030 | 031 | 032 | 033 | 034 | 035 | 036 | 037 | 038 | 039 | 040 | 041 | 042 | 043 | 044 | 045 | 046 | 047 | 048 | 049 |
| 2 | 049 | 050 | 051 | 052 | 053 | 054 | 055 | 056 | 057 | 058 | 059 | 060 | 061 | 062 | 063 | 064 | 065 | 066 | 067 | 068 | 069 | 070 | 071 | 072 | 073 | 074 | 075 | 076 | 077 | 078 | 079 | 080 | 081 | 082 | 083 | 084 | 085 | 086 | 087 | 088 | 089 | 090 | 091 | 092 | 093 | 094 | 095 | 096 | 097 |
| 3 | 097 | 098 | 099 | 100 | 101 | 102 | 103 | 104 | 105 | 106 | 107 | 108 | 109 | 110 | 111 | 112 | 113 | 114 | 115 | 116 | 117 | 118 | 119 | 120 | 121 | 122 | 123 | 124 | 125 | 126 | 127 | 128 | 129 | 130 | 131 | 132 | 133 | 134 | 135 | 136 | 137 | 138 | 139 | 140 | 141 | 142 | 143 | 144 | 145 |
| 4 | 146 | 147 | 148 | 149 | 150 | 151 | 152 | 153 | 154 | | | 156 | 157 | 158 | 159 | 160 | | | 161 | 162 | 163 | 164 | 165 | 166 | 167 | | | 168 | 169 | 170 | | | | | | | | | | | | | | | | 171 | | 172 | 173 |
| 5 | 174 | 175 | 176 | 177 | 178 | 179 | 180 | | 181 | 182 | | | | 183 | 184 | 185 | 186 | 187 | | 188 | 189 | 190 | 191 | | | | 192 | | 193 | | | | | | | | | | | | | | | | | | 194 | 195 | | 196 | 197 |
| 6 | 198 | | 199 | 200 | 201 | | | | 202 | 203 | | | | 204 | 205 | 206 | 207 | | | 208 | 209 | 210 | 211 | | | | 212 | 213 | | | | | | | | | | | | | | | | | | | 214 | 215 |
| 7 | 216 | | 217 | 218 | | | | | 219 | 220 | | | | 221 | | 222 | 223 | | | 224 | 225 | 226 | | | | | 227 | 228 | | | | | | | | | | | | | | | | | | | | 229 |
| 8 | 230 | | 231 | | | | | | 232 | 233 | | | | 234 | 235 | | | | | 236 | 237 | 238 | | | | | | | | | | | | | | | | | | | | | | | | | | 239 |
| 9 | 240 | | 241 | | | | | | 242 | 243 | | | | 244 | 245 | | | | | 246 | 247 | 248 | | | | | | | | | | | | | | | | | | | | | | | | | | 249 |
| 10 | 250 | | 251 | | | | | | 252 | | | | | 253 | 254 | | | | | 255 | | | | | | | | | | | | | | | | | | | | | | | | | | | | 256 |
| 11 | 257 | | 258 | | | | | | | | | | | 259 | | | | | | 260 | | | | | | | | | | | | | | | | | | | | | | | | | | | | 261 |
| 12 | | | 262 | | | | | | | | | | | | | | | | | 263 | | | | | | | | | | | | | | | | | | | | | | | | | | | | 264 |
| 13 | | | 265 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | 266 |
| 14 | | | 267 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | 268 |
| 15 | | | 269 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | 270 |
| 16 | 271 | 272 | 273 | 274 | 275 | 276 | 277 | 278 | 279 | 280 | 281 | 282 | 283 | 284 | 285 | 286 | 287 | 288 | 289 | 290 | 291 | 292 | 293 | 294 | 295 | 296 | 297 | 298 | 299 | 300 | 301 | 302 | 303 | 304 | 305 | 306 | 307 | 308 | 309 | 310 | 311 | 312 | 313 | 314 | 315 | 316 | 317 | 318 | 319 |
| 17 | 319 | 320 | 321 | 322 | 323 | 324 | 325 | 326 | 327 | 328 | 329 | 330 | 331 | 332 | 333 | 334 | 335 | 336 | 337 | 338 | 339 | 340 | 341 | 342 | 343 | 344 | 345 | 346 | 347 | 348 | 349 | 350 | 351 | 352 | 353 | 354 | 355 | 356 | 357 | 358 | 359 | 360 | 361 | 362 | 363 | 364 | 365 | 366 | 367 |
| 18 | 368 | 369 | 370 | | 371 | | | | 372 | | | | | 373 | 374 | 375 | | 376 | 377 | | | 378 | | | | | | | | | | | | | | | | | 379 | 380 | | | | 381 | | | | 382 | 383 |
| 19 | 387 | | | | 388 | | | | | | | | | 389 | | | | | | | | | | | | | | | | | | | | | | | | | 390 | | | | | | | | | | 391 |
| 20 | 392 | | | | | | | | | | | | | | | | | | | 393 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | 394 |
| +/- | 8 | 2 | 3 | 4 | 12 | 2 | 2 | 6 | 7 | 0 | 1 | 4 | 3 | 7 | 8 | 2 | 0 | 6 | 9 | 3 | 1 | 2 | 0 | 2 | 0 | 0 | 0 | 4 | 4 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 3 | 12 |
| +/- | 3 | 1 | 1 | 0 | 0 | 1 | 0 | 2 | 0 | 0 | 0 | 1 | 0 | 2 | 1 | 0 | 1 | 3 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 1 | 1 | 0 | 0 | 0 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 0 | 3 | 1 |

510 480 353 623 663 937

There is one serious issue that can be exploited for a cipher-breaking attack in Sniph that I am aware of. Regarding plaintext attacks, a single paired plaintext and ciphertext cannot meaningfully produce the correct keys, however, if multiple plaintexts are given with their matching ciphertexts, I believe there may only be *some* keysets that match up in the same dimensionalities between the pairs. The more ciphertexts you test will narrow down, very swiftly, which of those keysets is valid because very few, if more than one at all, will work with the same dimensionality each time.

To combat this, Sniph's code will be augmented soon so that each transformation process also includes unique charset sizes. Of the steps in Sniph so far, the Caesar shift and table dimensions are easy to understand, but the *charset size* may not be as easy to understand. Each plaintext character will get its own charset generated for it as well, with the minimum size being 96 (since 96 base characters need to be represented), but *N* additional characters repeated at the end. This is not the same as shifting the character set, i.e. 3 additional characters means that the character set will wrap as 'XYZABCABCDEF' and so on, creating a stutter in the charset. A separation of the passphrase will be used to determine the stutter, the reason for this being that it forces a separation analogous to the table dimensions being determined separately from the Caesar shift.

Since the passphrase as a whole will still determine the variables after these changes, and the process for their generation will be known, if there is a truly devastating flaw with library surfing, then it will be exposed even after this addition.

In any circumstance, I hope we finally attain privacy.

---