

A

We need to check if the first 9 or last 9 characters of s are equal to the given pattern.

You can use list slicing in Python, or `s.substr` in C++ in order to easily grab a contiguous substring of some string s .

```
1 s = input()
2 pattern = '...---...'
3
4 print("SOS" if pattern == s[:len(pattern)] or pattern == s[-len(pattern):] else "Hay SOS!")
5
```

B

We need to check if the given pattern appears as a *contiguous substring* anywhere in s . In Python, this is easily accomplished with the `in` operation. In C++, you can use `s.find`

```
1 s = input()
2 pattern = '...---...'
3
4 print("SOS" if pattern in s else "Hay SOS!")
5
```

But how does this work behind the scenes? What is its running time? For educational purposes, let's peek behind the curtain.

Let $|w|$ denote the length of some string w . For convenience, let $n := |s|$.

The string s has $|s| - |\text{pattern}| + 1$ contiguous substrings that are the same length as the given pattern (think of how many valid *starting indices* there are). Let's just check all of them and see if any are equal to our pattern.

```
1 s = input()
2 pattern = '...---...'
3
4 def is_substring(s, pattern):
5     for i in range(len(s) - len(pattern) + 1):
6         if s[i : i+len(pattern)] == pattern:
7             return True
8     return False
9
10
11 print("SOS" if is_substring(s, pattern) else "Hay SOS!")
```

Is string comparison free? Well, it's not a primitive operation. For educational purposes, let's unpack that as well. How does string comparison work?

- If the two strings are of different lengths then they are definitely not equal.
- Otherwise, the two strings are equal if and only if they match at each position.

You can imagine that if the strings are long, then checking that they match at *all* positions can be expensive, because we have to check *all* the indices. At worst, we end up checking almost all of the indices, which means the running time is $\mathcal{O}(|\text{pattern}|)$.

```
1 s = input()
2 pattern = '...---...'
3
4 def string_eq(a, b):
5     if len(a) != len(b):
6         return False
7     for i in range(len(a)):
8         if a[i] != b[i]:
9             return False
10    return True
11
12 def is_substring(s, pattern):
```

```

13     for i in range(len(s) - len(pattern) + 1):
14         if string_eq(s[i : i+len(pattern)], pattern):
15             return True
16     return False
17
18
19 print("SOS" if is_substring(s, pattern) else "Hay SOS!")

```

This is what's really happening under the hood when you use `in` or `find` between two strings. Thus, the running time of our solution is $\mathcal{O}(|\text{pattern}|(n - |\text{pattern}|))^1$. Since the pattern of `...---...` in our problem is only length 9, then the running time of our solution is $\mathcal{O}(9(n - 9))$, which we simplify to $\mathcal{O}(n)$, which is acceptable.

However, in general, it may be possible that $|\text{pattern}| \approx n/2$, which would cause our running time to degenerate to $\mathcal{O}\left(\frac{n}{2}(n - n/2)\right)$, which is $\mathcal{O}(n^2/4)$, which simplifies to $\mathcal{O}(n^2)$, which is too slow! In these cases, we would need to use a more advance technique like hash functions and rolling hashes. But that's overkill for this simple problem, so we won't discuss that any further here.

¹We can drop the +1 inside because remember, Big O Notation is just a rough estimate for large input sizes

K


We want to check if the given pattern `...---...` appears as a not-necessarily-contiguous *subsequence* in *s*.

A solution for this special pattern

I claim that `...---...` appears as a subsequence in *s* if and only if there are exactly three dashes between the first three dots and the last three dots.

```
1 s = input()
2 n = len(s)
3
4 prefix_dots = 0
5 first_third_dot_at = None
6 for i in range(n):
7     if s[i] == '.':
8         prefix_dots += 1
9         if prefix_dots == 3:
10             first_third_dot_at = i
11             break
12
13 suffix_dots = 0
14 last_third_dot_at = None
15 for i in range(n-1, -1, -1):
16     if s[i] == '.':
17         suffix_dots += 1
18         if suffix_dots == 3:
19             last_third_dot_at = i
20             break
21
22 between_dashes = 0
23 if first_third_dot_at is not None:
24     for i in range(first_third_dot_at+1, last_third_dot_at):
25         if s[i] == '-':
26             between_dashes += 1
27
28 print("SOS" if between_dashes == 3 else "Hay SOS!")
29
```

After all, when doing a subsequence check, *why not* just always take the earliest three dots in *s* as the first three dots of the pattern? If I take earlier dots, then there's "more of *s* left over", which makes it more likely that the rest of the pattern is contained in *s*. Having "more of *s*" is never worse!



The same argument applies for always taking the last three dots.


We only have to loop over the characters of s at most 3 times. Therefore, our solution runs in $\mathcal{O}(n)$.

A general solution for subsequence checks

This solution works for any s and any pattern, even with arbitrary letters and characters.

Remember when we said *why not* just always take the earliest three dots, because it's never worse? Think about it... doesn't that just always apply, in general?

Consider the first letter of the pattern. Find its earliest occurrence in s . *Why not* just always take this? It doesn't benefit us to wait, because *having more of s is always better*.



Then, repeat—for the next letter of the pattern, find its earliest occurrence in the remaining part of s ... and so on. Either we exhaust all the letters of the pattern (in which case yes, the pattern is a subsequence of s) or we exhaust all the letter of s first (in which case no, the pattern is not a subsequence of s).

```
1 s = input()
2 pattern = '...---...'
3
4 j = 0
5 for i in range(len(s)):
6     if s[i] == pattern[j]:
7         j += 1
8         if j == len(pattern):
9             break
10
11 print("SOS" if j == len(pattern) else "Hay SOS!")
```

Since there is only one loop, it's clear that the running time is $\mathcal{O}(n + |\text{pattern}|)$.

D

First of all, we would like to begin by reminding you about the magic number: as a coarse rule of thumb, a Python program can run somewhere on the order of magnitude of $\approx 10^7$ operations per second, and a C++ can run an order of magnitude of $\approx 10^8$ operations per second. If we want our code to run under the time limit, then we need to make sure its operation count is roughly under this magic number.

We want to count how many integers satisfy both of these conditions:

- It lies between L and R (inclusive)
- It's an SOS number.

The most straightforward idea might look something like this: For each number from L to R , test each one for being an SOS number. See the following pseudocode.

```
1 total = 0
2 for each integer k from L to R:
3     if is_SOS(k):
4         total += 1
5 print(total)
```

There are $R - L + 1$ numbers to check, but in the worst case, $L = 1$ or some other small number, so we just say that there are $\mathcal{O}(R)$ candidate numbers. For each one, its check can take up to $\mathcal{O}(\text{no. of digits})$ operations (and we can show that this is $\mathcal{O}(\log_{10}(R))$). Therefore, the running time is $\mathcal{O}(R \log_{10} R)$. Unfortunately, with R possibly up to 10^{15} , this is *way* too slow.

The problem is that there are just too many integers from L to R . And if we look at the sample input, there aren't even that many SOS numbers from 1 to 10^{15} in the first place! Most of our checking just goes to waste.

"The integers from L to R " makes up our **search space**, because we are *searching* for SOS numbers in this range. The problem is that our search space is too large and too sparse. Is there maybe a better set that we can search through?

Well... how about we *swap* the search space and the check? See this pseudocode.

```
1 total = 0
2 for each SOS number k that is <= 1e15:
3     if L <= k <= R:
4         total += 1
5 print(total)
```

How large is this new search space? Well, SOS numbers $\leq 10^{15}$ must be strings of length at most 15 digits, each of which is either 5 or 0. Actually, because of how leading zeroes work with our number system, we can just generate *all* such strings with exactly 15 digits. Because each digit can be one of two possible values, there are exactly $2^{15} = 32768$ such strings, which is not too large.

Of course, not all these strings are SOS numbers, but that's fine! Just filter out only the ones with 505 as a substring. There's still wasted work, but at least 2^{15} is a tolerable amount of work.

We can enumerate all SOS numbers using recursive backtracking.

```
1 all_SOS_numbers = []
2 def generate_SOS_numbers(i, SOS):
3     if i == 0:
4         if '505' in SOS:
```

```

5         all_SOS_numbers.append(int(SOS))
6     else:
7         generate_SOS_numbers(i-1, SOS + '5')
8         generate_SOS_numbers(i-1, SOS + '0')
9
10 B = 15
11 generate_SOS_numbers(B, '')

```

As a special hack, since we only care about binary strings, we can use the **binary representations** of the integers from 0 to $2^{15} - 1$ in order to generate all 15-length patterns with two symbols. This technique is called *bitmasking* and has loads of other applications as well.

The following code uses the "bitwise operations" left shift and bitwise AND, which you can look up. Many languages also may have their own built-in ways to extract the binary representation of an integer. For Python, you can use `bin`, and for C++, you can use `std::bitset`.

```

1 all_SOS_numbers = []
2
3 B = 15
4 for mask in range(1 << B):
5     SOS = ''
6     for i in range(B):
7         if mask & (1 << i): # if ith bit is 1
8             SOS += '5'
9         else:
10            SOS += '0'
11
12 if '505' in SOS:
13     all_SOS_numbers.append(int(SOS))

```

In summary, this is our final code:

```

1 all_SOS_numbers = []
2
3 ... # choose your preferred method of generating the SOS numbers
4
5 L, R = map(int, input().split())
6 total = 0
7 for SOS in all_SOS_numbers:
8     if L <= SOS <= R:
9         total += 1
10 print(total)

```

Let $B = \log_{10}(R)$ be the number of digits in the largest SOS number to consider. Then, our solution runs in $\mathcal{O}(B2^B)$, because there are 2^B two-character strings of length B , and testing each one for SOS-ness takes $\mathcal{O}(B)$ time, which is fast enough.

Bonus: The recursive backtracking solution can be modified so that the solution runs purely in $\mathcal{O}(2^B)$. Not that it's too important, because the bounds are small enough, but... can you think of how?

Bonus for experienced participants: Suppose there are T different test cases, each with their own L and R and asking how many SOS numbers lie in that range. Here, $1 \leq T \leq 2 \times 10^5$, so T can be quite large. Can you think of a fast way to answer each of these test cases? Say, suppose you should be able to answer each test case in $\mathcal{O}(\log(R))$ time. *Hint:* You can do some pre-processing before answering all the test cases.

Bonus Bonus for the *very* experienced participants: Suppose that $1 \leq L \leq R \leq 10^{2 \times 10^5}$, i.e. L and R can have up

to 2×10^5 *digits*. As is standard in comp prog, suppose we only want the answer modulo $10^9 + 7$. Can we answer this in $\mathcal{O}(B)$ (not $\mathcal{O}(2^B)$)?