

CS358: Applied Algorithms

Assignment 3: Cuckoo Filters (due 4/16/2020)

Instructor: Sam McCauley

Instructions

All submissions are to be done through github. This process is detailed in the handout “Handing In Assignments” on the course website. Answers to the questions below should be submitted by editing this document. All places where you are expected to fill in solution are marked in comments with “FILL IN.”

Please contact me at sam@cs.williams.edu if you have any questions or find any problems with the assignment materials.

Problem Description

Let’s call a word a “bilingual palindrome” if the reverse of that word is a word in another language. For example, “mures” is a bilingual palindrome, because “mures” is a word in French, and “serum” is a word in English. A palindrome may not be a bilingual palindrome: for example, “racecar” is not a bilingual palindrome unless “racecar” is a word in another language.

Given a list of words in several languages, we can efficiently find all the bilingual palindromes by storing all the words in a hash table (along with what languages they appear in), and then going through each word and looking up its reverse in the hash table. If the reverse appears in a language different from the original word, then each is a bilingual palindrome.

This assignment asks you to implement a cuckoo filter to speed up this lookup process. Before checking the whole hash table for the reverse, we first check the cuckoo filter. We only check the hash table if the cuckoo filter says that the reverse may be present. Even though we are dealing relatively small amounts of data (the original dictionaries have total size less than 10MB), a cuckoo filter speeds up the implementation by a factor of approximately 2! This effect would become much more significant for larger datasets.

The only coding requirement for this assignment is to implement a cuckoo filter. The hash table, searching, output, etc., are all already implemented in `test.c`. You may change `test.c` if you wish, but you’ll only be graded on your cuckoo filter implementation.

The full problem description is given here to help with your understanding.

INPUT: `test.out` is given five arguments. The first three are strings representing word lists. The fourth is an integer representing the total number of (unique) words across these lists. The fifth is an integer representing the number of bilingual palindromes in these lists (that is to say, the fifth number is the desired answer).

Each word list is assumed to consist of a sequence of words, one on each line. Blank lines are ignored. The words are assumed to not have punctuation.

Note that these word lists are stored in ASCII format. In particular, this means that accented characters¹ were (automatically and likely poorly) transliterated into English equivalents—so people familiar with some of these languages may notice that things are slightly off.²

The following command runs your program with the correct number of inserted words and expected answer:

```
./test.out english.txt french.txt german.txt 916641 1977
```

You may also test your command with the much shorter `dictionaryShort.txt`.

OUTPUT: The testing program automatically checks if the number of bilingual palindromes found matches the desired number. In this assignment, the program automatically gives some extra output that may be helpful. Further output (i.e. listing the palindromes) can be toggled using `#define` statements.

FILTER FUNCTIONS: I have given you a starting point for how to use the filter—the testing program already interfaces with the following filter functions. I’ve also given the outline of a `struct` which I found useful to store some of the filter metadata; it can be found in `filter.h`. It’s quite possible that you will change this struct, or you may not even want to use it at all. You may also edit the functions any way you wish—i.e. changing the type of the arguments or how many there are (provided that you are still implementing cuckoo filter functionality).

To be clear, you can get a perfect grade on the coding portion of this assignment while only editing the inside of the functions (denoted by comments) in `filter.c`.

You may notice that there is a (global) array at the beginning of `filter.c`. Its entries are random. The intention is to use this array as the hash h which you apply to the fingerprint in partial-key cuckoo hashing.

Here is a list of the functions and how they are used:

```
void filter_instantiate(Filter* filter, int numWords)
```

This function is called before any other calls to the filter. You can think of it like a constructor. It should set constants and allocate memory. `filter` is a pointer to a struct that I found useful (you can change this to not use a struct if you wish); `numWords` is a guarantee on how many elements will ever be inserted into the cuckoo filter. (In other words, `numWords` is n .)

```
void filter_insert(char* word, int length, Filter* filter)
```

This function inserts a new word (given by `word`) into the filter. `length` is the length of the word, and `filter` is a pointer to the filter we want to insert to. After `word` is inserted, if we call `filter_lookup` on `word` and `filter` it should always return 1.³

¹Characters like é or ü or ß; really anything other than A-Z.

²I put a decent amount of effort into trying to get around this but C really really does not play well with non-ASCII characters.

³You can test this with `CHECK_CORRECTNESS` in `test.c`.

`test.c` will insert each *unique*⁴ word in each word list into the filter by calling this function. These inserts will all occur before any call to `filter_lookup`.

```
int filter_lookup(char* word, int length, Filter* filter)
```

This function looks up a word (given by `word`) in the filter—it is equivalent to the “is $q \in S$?” query discussed in class. `length` is the length of the word, and `filter` is a pointer to the filter we want to insert to. This function returns 1 to represent the filter saying “ $q \in S$ ”, and 0 to represent the filter saying “ $q \notin S$ ”. Note that, in the starter code, this function always returns 1, thus satisfying the No False Negatives guarantee.

`test.c` will call this function before querying the hash table for each reversed word; it will only query the hash table if this function returns 1.

PARAMETERS: Your cuckoo filter should implement the following parameters. Each fingerprint should be of length 8 bits. There should be 5%-10% space overhead; that is to say, the number of total slots should be around $1.10 * (\text{num_words})$.⁵ You should use partial-key cuckoo hashing, so you should have two hash functions h_1 and $h_2(x) = h_1(x) \wedge h(f(x))$. Each hash entry should consist of 4 slots, where each slot is large enough to store a fingerprint.⁶

On an insert, your filter should exit if the number of “cuckoo” iterations exceeds `max_iter`. I would suggest setting `max_iter` = $4 \log_2(\text{num_words})$. If this causes frequent failures, feel free to make this number higher.

Questions

Problem 1 (65 points). Your code is worth 65 points. **You are not required to write about your implementation to get credit for this assignment; the 65 points are for code alone.** But, you may fill in some information below if there is something you want me to know about your implementation.

Solution.

Problem 2 (10 points). Give an upper bound on the false positive rate ε of your filter. That is to say, analyze the (theoretical) false positive probability with bins of size 4, fingerprints of length 8, 1.05 fingerprint slots per element, and 2 hash functions.

Hint: Most of the work for this question was done in the lecture.

Solution.

Problem 3 (10 points). Let’s say an element we’re inserting is *easy* if it can be inserted immediately, without cuckooing. Give an upper bound on the probability that an element is easy.

Hint: An element is easy if the first slot we look at is empty, OR the second slot is empty, OR the third slot is empty, and so on.

⁴You do not need to worry about duplicates; they will be removed before this function is called.

⁵1.05 works fine for me, but it’s OK if a little more space helps make the filter work.

⁶This conveniently means that each bin is of size 32 bits.

Solution.

Problem 4 (15 points). Let's say I take a Bloom filter for a set S_1 using a bit array A_1 and hash functions h_1, \dots, h_k , and a Bloom filter for a set S_2 using a bit array A_2 and the same hash functions h_1, \dots, h_k .

I create a bit array T using the bitwise OR of A_1 and A_2 . (That is to say, I get a bit array T where $T[i] = A_1[i] \vee A_2[i]$.) Then T along with h_1, \dots, h_k is a Bloom filter—it is a bit array with an associated hash function.

(a) If I query an element q , and its bits are all set to 1, (i.e. if $T[h_i(q)] = 1$ for all i), what can I say about q ? (For example: can we say that $q \in S_1$ for sure? Or that q is unlikely to be in S_2 ?)

(b) If I query an element q , and one of its bits is 0 (i.e. $T[h_i(q)] = 0$ for some i), what can I say about q ?

(c) Assume that $k = \log_2(1/\varepsilon)$ and $|A_1| = |A_2| = 1.44n \log_2(1/\varepsilon)$ (where $n = |S_1| = |S_2|$). As we saw in class, we can analyze these filters by saying that each bit of A_1 and A_2 is 1 with probability $1/2$. Under this assumption, what is the probability that each bit of T is 1? (I am looking for an exact answer, not an upper bound.)

(d) (**Extra credit: 5 pts**) Using your answer to part (c), what is the false positive rate of T ?

Solution.