

15-122: Principles of Imperative Computation

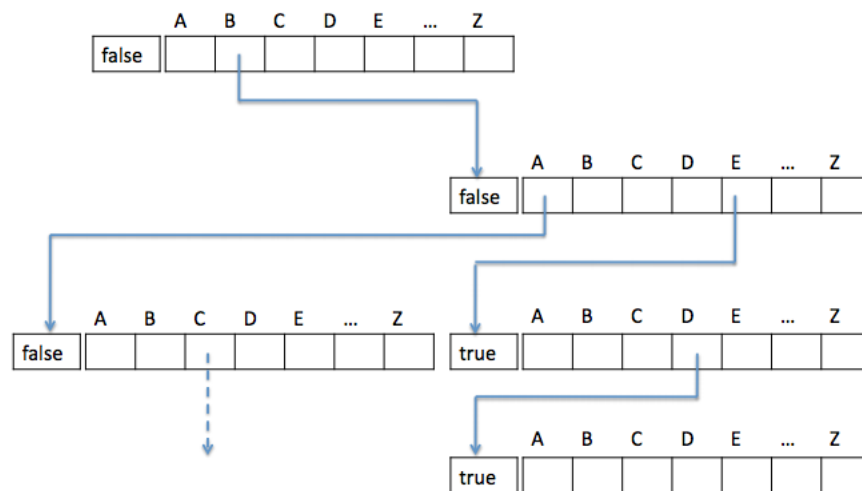
Recitation 22

Matt Dee, Josh Zimmerman

Trie-ing is Cool!

The basic idea of a trie is that, for some purposes, like spell checking and boggle, it often makes sense to see if a string is the *prefix* of some word in a dictionary. (That way, we can realize that we can stop checking some sequence of letters in boggle, and more quickly flag that a word is misspelled, without checking every word)

To do this we keep a tree that contains letters representing prefixes. Each sub-trie is the of all strings in the original trie that started with the letter of the parent. The diagram below shows an example of this.(The diagram is an example of a *multi-way* trie, which is one implementation of the basic idea of looking up prefixes of strings.)



This multi-way trie contains the words “be”, “bed”, and the start of “baccalaureate”. Image credit: Frank Pfenning.

The boolean flags in the trie are true if and only if that node in the trie is the end of a word. (It’s useful to know when we’ve reached the end of a word so we can know that, for instance, “b” and “bacca” are not words, but are prefixes for other words)

This approach is great, since it lets us look up whether a word of length k is in the trie in $O(k)$ time (the size of the dictionary does not matter for runtime!).

However, it has some major problems. In particular, we’ll be wasting a *lot* of space, since we have an array of 26 elements at each level, of which we’ll only need a few the majority of the time. And that’s just for lowercase! If we tried to make a trie for all ASCII characters (128 chars), we would be in real trouble, and Unicode is out of the question (>230,000 chars)!

We don’t have unlimited memory, so this is bad.

So, we’ll use a slightly more complicated but far smaller structure, whose size does not depend so much on the size of the dictionary, a *ternary search trie*.

Implementation

We will represent these using the following structs, using a structure very reminiscent of a BST:

```
1 typedef struct tst_node tnode;
2 struct tst_node {
3     char c; /* discriminating character */
4     elem data; /* possible data element (void pointer) */
5     tnode *left; /* left child in bst */
6     tnode *middle; /* subtrie following c */
7     tnode *right; /* right child in bst */
8 };
9 typedef struct trie_header trie
10 struct trie_header {
11     tnode *root;
12 };
```

Probably the most important functions of the trie are the `trie_lookup` functions and `trie_lookup_prefix`, so we will look at the code to understand how lookups work, and hopefully get a better idea of how data in tries is organized. Both of these functions share some logic, so we will write a helper function to get the `tnode` for a string.

```
1 /* Returns the tnode for the given string, starting at character i */
2 tnode *tnode_lookup(tnode *T, char *s, size_t i)
3 {
4     REQUIRES(is_tnode_root(T));
5     REQUIRES(s != NULL);
6     REQUIRES(i < strlen(s));
7     if (T == NULL) return NULL;
8     // If the character we are looking for is less than this node, look left
9     if (s[i] < T->c) return tnode_lookup(T->left, s, i);
10    // if the character is greater, look right
11    if (s[i] > T->c) return tnode_lookup(T->right, s, i);
12    ASSERT(s[i] == T->c);
13    // We have found the character!
14    if (s[i+1] == '\0') return T; //We are at the end, return this node
15    return tnode_lookup(T->middle, s, i+1); //Continue the search on the next character
16 }
17 /* Returns a tnode* if it is a valid prefix, NULL otherwise */
18 tnode *trie_lookup_prefix(trie *TR, char *s) {
19     REQUIRES(is_trie(TR));
20     REQUIRES(s != NULL);
21     return tnode_lookup(TR->root, s, 0);
22 }
23 /* Returns the elem associated with s, if it exists */
24 elem trie_lookup(trie TR, char *s)
25 {
26     REQUIRES(is_trie(TR));
27     REQUIRES(s != NULL);
28     tnode *T = trie_lookup_prefix(TR, s);
29     if (T == NULL) // This isn't even a valid prefix...
30         return NULL;
31     else // Return the element there (may be NULL if s is a prefix but not a key)
32         return tnode_elem(T);
33 }
```

Invariants

Tries, like any good data structure, have invariants that we need to ensure hold true. Specifically, let's look at the `is_tnode` function for ternary search tries. Because `NULL` represents the empty trie, this is a valid `tnode`. For any non-null trie, the trie invariants are:

- The `left` and `right` fields act as a BST. In other words, for any `tnode T`, IGNORING `middle` pointers, every `tnode` in the `T->left` subtree represents a character whose ASCII value is less than `T->c`, and every `tnode` in the `T->right` subtree represents a character whose ASCII value is greater than `T->c`.
- For any `tnode T`, `T->middle` is a valid `tnode`.
- If, for some `tnode T`, `T->middle == NULL`, then `T->elem` must not be `NULL` (otherwise this trie would not be a prefix for any element!).

To check this, we use mutually recursive functions. This refers to two functions, which each call each other. Note that in order to do this, we have to forward-declare the second function so that the first function won't cause a compiler error when it calls the second.

```
1 // We need this declaration so that is_tnode knows about is_tnode_root
2 bool is_tnode(tnode *T, int lower, int upper);
3 bool is_tnode_root(tnode *T);
4
5 // Ensure that, each node is a valid tnode, and that every character on this level
6 // (without taking a middle) is less than lower and greater than upper
7 bool is_tnode(tnode *T, int lower, int upper) {
8     if (T == NULL) return true; // Empty tnode is valid tnode
9     // Check bounds (BST invariant)
10    if (!(lower < T->c && T->c < upper)) return false;
11    // Recursively check children
12    if (!(is_tnode(T->left, lower, T->c))) return false;
13    if (!(is_tnode(T->right, T->c, upper))) return false;
14    // T->middle does not need to be within the same limits since it is a "root"
15    // for another tree so we call this to check those invariants.
16    if (!(is_tnode_root(T->middle))) return false;
17    if (T->middle == NULL && T->data == NULL) return false;
18    return true;
19 }
20 bool is_tnode_root(tnode *T) {
21     return is_tnode(T, 0, ((int)CHAR_MAX)+1);
22 }
```

Note that `is_tnode_root` is the entry point, checking validity for a root node (whose character value has no limits), and `is_tnode` checks the ordering invariants that come with the BST structure in addition to checking that each `middle` field points to a valid root node.

Can we trie Boggle?

Boggle is a word game that we can use a trie to solve. The basic idea is that we have a four-by-four grid of random letters. We try to make words out of them by drawing a line through letters. We can use the trie interface to help us with this. Think about how, and how we might improve the interface to allow us to find all words in a boggle board as quickly as possible.