

# Exam 2

April 19th, 2012

Work alone. Do not use any notes or books. You have approximately 75 minutes to complete this exam.

Please write your answers on the exam. More paper is available if you need it. Please put your name at the top of the first page.

There are four questions on this exam, for a total of 80 points.

## 1 Pointer juggling (20 points)

On which line will the following program fail with a segmentation fault, and why?

```
1 #define N (3)
2
3 int
4 main(int argc, char **argv)
5 {
6     int i;
7     struct pointy {
8         struct pointy *p[2];
9     } a[N];
10
11     for(i = 0; i < N; i++) {
12         a[i].p[0] = &a[i];
13         a[i].p[1] = a + (i + 1) % N;
14     }
15
16     a[0].p[0] = 0;
17
18     a[2].p[0]->p[1]->p[1] = a[0].p[0];
19     a[1].p[0]->p[0]->p[1] = a[0].p[0];
20     a[0].p[1]->p[1]->p[0] = a[2].p[0];
21     a[0].p[0]->p[1]->p[1] = a[2].p[1];
22     a[0].p[0]->p[1]->p[1] = a[2].p[0];
23
24     return 0;
25 }
```

### Solution

We get a segmentation fault on line 20, when the program attempts to dereference `a[0].p[1]`. The reason is that `a[0].p[1]` is assigned a null pointer in line 18, since `a[2].p[0]->p[1]` is `&a0`. This is not affected by line 19, which puts a null pointer in `a[1].p[1]`.

## 2 Choosing a data structure (20 points)

Suppose you are asked to write a program for displaying a leaderboard for a large on-line multiplayer game. The leaderboard consists of the names of the 10 players with the highest scores and their scores. Each player of the game has a name (a string) and a score, equal to the number of games they have won. Your program should provide routines `incrementScore(const char *player)` and `returnWinners(char *winners[10], int scores[10])` where `incrementScore` increases the score of the named player by 1 (or sets the score to 1 if the player has never won a game before) and `returnWinners` fills in the `winners` and `scores` arrays with the names and scores of the players with the 10 highest scores (with ties broken arbitrarily).

Without writing actual code, describe as succinctly as you can what data structure would be a good choice for this task and why.

### Solution

The simplest solution would probably be to keep track of the scores for all players using a hash table and keep track of the top 10 players in a sorted array. The reason this works is that a new player can only enter the top 10 by displacing the lowest existing player. So when a score is updated, we can (a) update the score in the hash table; (b) check using a linear scan if the player is already on the leaderboard, and float them up if their score has increased enough to warrant it; or (c) if the player is not on the leaderboard but has a score greater than the lowest leader's score, replace the lowest leader and float up if needed.

Since the size of the leaderboard is small and constant, it would also work to have an unsorted leaderboard, but keeping it sorted saves time in `returnWinners` and in finding the smallest score.

Another reasonable option would be a priority queue for the leaderboard indexed by a hash table or some other plausible dictionary data structure for the player names. The downside of this approach compared to the simple array is that it adds complexity, increases the cost from  $O(1)$  to  $O(\log n)$ , and requires more space to store the non-leaders in the priority queue.

### 3 A digital search tree (20 points)

Here is an implementation of the creation and insertion routines for a binary digital search tree. The idea of the tree is that a key of type `unsigned int` is represented by a path through the tree, where each edge in the path is determined by the next bit in the key.

Write a routine `DSTpredecessor` that returns the largest key in the tree less than or equal to the given key, or 0 if there is no such key. We have provided a declaration for this function for you on the next page.

```
#include <stdlib.h>
#include <limits.h>

#define BITS (sizeof(unsigned int) * CHAR_BIT) /* CHAR_BIT == 8 */

struct node {
    struct node *kids[2];
};

struct node *DSTcreate(void)
{
    struct node *root;

    root = malloc(sizeof(struct node));
    root->kids[0] = root->kids[1] = 0;
    return root;
}

void DSTinsert(struct node *root, unsigned int key)
{
    unsigned int mask;
    int bit;
    for(mask = 1 << (BITS-1); mask != 0; mask >>= 1) {
        bit = (key & mask) != 0;
        if(root->kids[bit] == 0) {
            root->kids[bit] = malloc(sizeof(struct node));
        }
        root = root->kids[bit];
    }
}
```

```

/* return largest key <= key or 0 if there is none */
unsigned int DSTpredecessor(struct node *root, unsigned int key);

```

## Solution

Here's one possible solution, using an auxiliary function that finds the smallest value in a subtree:

```

/* return smallest key in subtree, not including bits above root */
static unsigned int
smallestKey(struct node *root)
{
    int ret = 0;

    for(;;) {
        if(root->kids[0]) {
            root = root->kids[0];
            ret = ret << 1;
        } else if(root->kids[1]) {
            root = root->kids[1];
            ret = (ret << 1) + 1;
        } else {
            return ret;
        }
    }
}

```

```

/* return largest key <= key or 0 if there is none */
unsigned int DSTpredecessor(struct node *root, unsigned int key)
{
    unsigned int mask;
    unsigned int bit;
    unsigned int prefix;    /* bits we've accumulated so far */

    prefix = 0;

    for(mask = 1 << (BITS-1); mask != 0; mask >>= 1) {
        bit = (key & mask);
        if(root->kids[1]
            && bit != 0
            && (prefix + mask + smallestKey(root->kids[1]) <= key)) {
            /* it's in the right subtree */
            prefix += mask;
            root = root->kids[1];
        } else if(root->kids[0]) {
            if(bit == 0) {
                /* it's in the left subtree */

```

```
        root = root->kids[0];
    } else {
        /* it's the biggest value in the left subtree */
        key = ~0;
        root = root->kids[0];
    }
} else {
    /* nowhere to go */
    return 0;
}
}

return prefix;
}
```

## 4 Change this program (20 points)

Here is a rather dangerous implementation of a queue. Turn it into a stack by changing exactly one line of code.

```
1 #include <stdlib.h>
2
3 struct thing {
4     int size;          /* size of contents */
5     int bot;           /* first used location */
6     int top;           /* first unused location */
7     int *contents;    /* array of elements */
8 };
9
10 struct thing *thingCreate(int size)
11 {
12     struct thing *t;
13
14     t = malloc(sizeof(*t));
15
16     t->size = size;
17     t->bot = t->top = 0;
18     t->contents = malloc(sizeof(int) * size);
19
20     return t;
21 }
22
23 void thingPush(struct thing *t, int value)
24 {
25     t->contents[t->top++] = value;
26 }
27
28 int thingPop(struct thing *t)
29 {
30     return t->contents[t->bot++];
31 }
```

### Solution

Change line 30 to

```
30     return t->contents[--(t->top)];
```