Collaboration: In lab, we encourage collaboration and discussion as you work through the problems. These activities, like recitation, are meant to get you to review what we've learned, look at problems from a different perspective and allow you to ask questions about topics you don't understand. We encourage discussing problems with your neighbors as you work through this lab!

Setup: Copy the lab code from our public directory to your private directory:

```
% cd private/15122
% mkdir lab13
% cd lab13
% wget https://web2.qatar.cmu.edu/~srazak/courses/15122-s20/lab/handout-13.tgz
% tar xfvz handout-13.tgz
```

Grading: Finish task (1.a) for full credit, and finish both tasks (1.a) and (1.b) for extra credit.

Generic sort

Today we will explore the difference between **void**∗ in C0 and **void**∗ in C, and exploit that difference to write a generic sort in C.

The image above shows a 16 byte allocation where each byte contains an 8-bit value. The actual size of a byte in C is implementation-defined and so can be 8 bits *or more*. You can pretty much count on a byte being 8 bits on current computers.

If we have a pointer **A** whose value is `0xB58`, then the way we interpret that pointer *depends on the type of the pointer*. As a **char**∗, the pointer **A** points to the value `0x48` or `'H'`, the first character in the NUL-terminated string `"Hello, everyone"`. As an `int32_t∗`, the pointer **A** points to a signed integer, the first element in an array of four integers. (According to the implementation-defined behavior we usually expect, this array contains the four integers 1819043144, 1696607343, 2037540214, and 6647407.) If **A** is a **void**∗, then we know *nothing* about how to interpret, read from, or write to this block of memory.

In this lab, we will write a function that sorts arrays without knowing anything about how to interpret, read from, or write to the memory addresses in that array. The client will tell us the how many elements there are (`count`) and the size of each element (`elt_size`). The client will also tell us how to interpret (with a comparison function) and manipulate (with a swap function) elements of the array.

```
typedef void swap_fn(void *x, void *y)
  /*@requires x != NULL && y != NULL; @*/ ;

// Compares the values at memory locations x and y
typedef int compare_fn(void *x, void *y)
  /*@requires x != NULL && y != NULL; @*/
  /*@ensures -1 <= \result && \result <= 1; @*/ ;

void gsort(void *A, size_t count, size_t elt_size,
           swap_fn *swp, compare_fn *cmp)
  /*@requires A != NULL && swp != NULL && cmp != NULL; @*/
  /* requires that A is an allocation of at least count * elt_size */ ;
```

The interface above is provided in `lib/gsort.h`. Your implementation in `gsort.c` is one of the very small number of cases in C where it is acceptable to cast a **void***∗ to another pointer type without being absolutely certain what type the **void**∗ originally was. (Remember that it was *never* acceptable to do this in C1.)

If a client asks us to sort the 5 two-byte values starting at the **void** pointer `0xB58`, then we know that the five elements in the array have the addresses `0xB58`, `0xB5A`, `0xB5C`, `0xB5E`, and `0xB60`. We can calculate the address of the array element that the client thinks of as `A[3]` by casting `A` to a **char**∗ and then writing either `(A + 6)` or `&A[6]`. This makes sense because a **char** is always one byte and because 6 is the array offset (3) times the size of an array element in bytes (2).

All we need to do to sort is calculate the addresses where array elements begin and pass these addresses to the client functions. The client's functions know how to compare and swap values given the addresses of those values. Because you are writing the sorting algorithm without knowing what values are being stored, you shouldn't ever access the memory in the array directly.

(1.a) Write a generic sort in `gsort.c` according to the strategy described above. Make sure your `gsort.c` includes `"lib/gsort.h"`. A C0 version of selection sort is given to you below and in `sort.c0`. Don't worry about translating the contracts. You can test your code by running:

```
% make
% ./a.out
```

Submit your code in Autolab to collect the points.

3pt

(1.b) Modify `freqinfo.c` to sort arrays of frequency counts by both frequency and by word (according to the **string**.h function `strcmp`).

You can use `gsort-test.c` to test your solution. To test, remove the return statement in between the two parts and add the testing code in the second part of the main function.

After that, you can test your code by running:

```
% make
% ./a.out
```

Submit your code in Autolab to collect the points.

4pt

Selection sort in C0

```c
1  void sort(int[] A, int n)
2  //@requires 0 <= n && n <= \length(A);
3  //@ensures is_sorted(A, 0, n);
4  {
5    for (int i = 0; i < n; i++)
6    //@loop_invariant 0 <= i && i <= n;
7    //@loop_invariant is_sorted(A, 0, i);
8    //@loop_invariant le_segs(A, 0, i, A, i, n);
9    {
10     int min = i;
11     for (int j = i + 1; j < n; j++)
12     //@loop_invariant i < j && j <= n;
13     //@loop_invariant i <= min && min < n;
14     //@loop_invariant le_seg(A[min], A, i, j);
15     {
16       if (A[j] < A[min])
17         min = j;
18     }
19     swap(A, i, min);
20   }
21 }
```