Instructions:

- Print your name in the space provided below.
- This examination is closed book and closed notes, aside from the permitted one-page formula sheet. No calculators or other computing devices may be used. The use of any such device will be interpreted as an indication that you are finished with the test and your test form will be collected immediately.
- Answer each question in the space provided. If you need to continue an answer onto the back of a page, clearly indicate that and label the continuation with the question number.
- If you want partial credit, justify your answers, even when justification is not explicitly required.
- There are 7 questions, some with multiple parts, priced as marked. The maximum score is 100.
- When you have completed the test, sign the pledge at the bottom of this page and turn in the test.
- If you brought a fact sheet to the test, write your name on it and turn it in with the test.
- Note that either failing to return this test, or discussing its content with a student who has not taken it is a violation of the Honor Code.

Do not start the test until instructed to do so!

Name ______**Solution**______________________________
printed

Pledge: On my honor, I have neither given nor received unauthorized aid on this examination.

____
signed

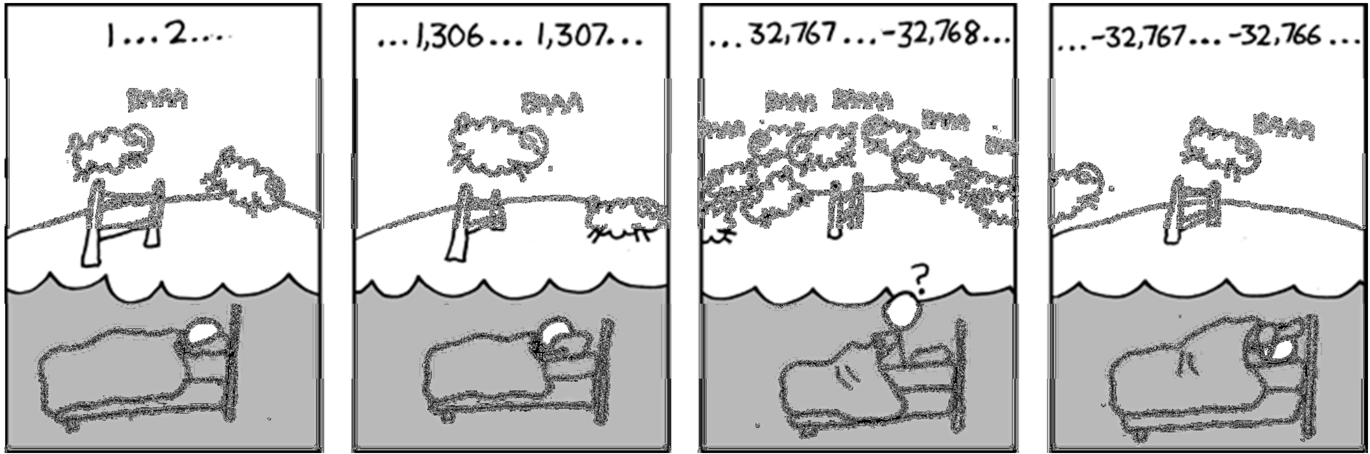

xkcd.com

1. [18 points] Perform the indicated conversions. You must show work in order to receive credit!

a) Convert from 8-bit 2's complement (i.e., signed) form to base-10: `1010 0101`

This is a negative value (high-order bit is 1). The corresponding positive value is 0101 1011 (flip the bits and add 1).

That represents 2^0 + 2^1 + 2^3 + 2^4 + 2^6 = 1 + 2 + 8 + 16 + 64 = 91.

So the original bit-sequence is equivalent to the base-10 representation -91.

b) Convert from 8-bit unsigned form to base-10: `1001 0111`

This is a straight base-2 representation, so it's 2^0 + 2^1 + 2^2 + 2^4 + 2^7 = 1 + 2 + 4 + 16 + 128 = 151.

c) Convert from base-10 to 8-bit unsigned binary: `29`

Successively divide by 2 and record remainders:

```
2  29
   14   1
    7   0
    3   1
    1   1
    0   1
```

So, in base-2 this is represented as 11101.

2. Recall that C supports both signed and unsigned integer types. For example, in C we have the signed type `int16_t` and the unsigned type `uint16_t`.

a) [8 points] What is the difference between the way an `int16_t` value and a `uint16_t` value would be represented in memory? Be precise.

Signed integer values, like `int16_t`, are represented in 2's complement.

Unsigned integer values, like `uint16_t`, are represented in pure base-2.

Both types would be represented using 16 bits.

There are other differences, like the difference in ranges, but those are consequences of the way the in-memory representation is done, differences in the in-memory representation.

b) [8 points] Briefly describe a programming situation in which having the `uint16_t` type available would lead to a better solution than if you only had the signed integer types available, and explain why.

The key difference is that `uint16_t` variables can store substantially larger values than `int16_t` variables; both types have the same size range, but the maximum for unsigned types is larger.

So, any scenario in which we need to represent only non-negative values, but the largest ones will exceed the maximum for an `int16_t` type, would illustrate the point.

One example would be in storing scores for a game (if the maximum likely score was no larger than 65000 or so).

Another example would be in numbering records for a small club or school; negative values would be irrelevant, and the upper limit for the `uint16_t` might be sufficient to handle records for a considerable length of time.

The most common issue here was in not giving a specific programming situation.

3. [10 points] Circle the strings below that would match the following regular expression:

```
(abc)?.[0-9]+(z|x|f)
```

a) `abcKf` b) `abcH8z` (circled) c) `%6x` (circled) d) `abcJ2t` e) `k00f` (circled)

a) doesn't contain a digit; d) doesn't end with a z, x or f.

---

4. [14 points] Formulate a `grep` regular expression that would match all, and only, those strings that satisfy all of the following conditions (strings that correspond to the sort of course labels VT uses):

- The string starts with 2, 3 or 4 upper-case letters.
- Those are followed immediately by 1, 2, or 3 spaces, corresponding to the number of letters in the first part, respectively.
- The remainder of the string consists of exactly 4 base-10 digits, the first of which is never 0.

So, for example, your regular expression should match "`CS•••2505`", "`SOC••1004`" and "`MATH•2534`", but not "`SOC•1004`", where '• ' denotes a space character.

1: 2, 3 or 4 upper-case letters: `[A-Z]{2,4}` **or** `[A-Z]{2}[A-Z]?[A-Z]?`

2: We have to tie the number of spaces to the number of characters in the first part:

```
([A-Z]{2}   |[A-Z]{3}  |[A-Z]{4} )
```

The RE contains the right number of spaces after each sequence of upper-case letters, but this won't work if there are extra spaces…

3: OK, this is easy: `[1-9][0-9]{3}`

Putting it all together:

```
([A-Z]{2}   [1-9]|[A-Z]{3}  [1-9]|[A-Z]{4} [1-9])[0-9]{3}
```

But… this does still contain a flaw; we need to specify that there's something different from an upper-case letter preceding all of that:

```
[^A-Z]([A-Z]{2}   [1-9]|[A-Z]{3}  [1-9]|[A-Z]{4} [1-9])[0-9]{3}
```

---

5. [12 points] Formulate a `grep` regular expression that would match all, and only, three-digit integers between 100 and 255. So, for example, it would match 243 and 105 but not 261.

`1[0-9]{2}` **matches 100 - 199**
`2[0-4][0-9]` **matches 200 - 249**
`25[0-5]` **matches 250-255**

So, this will do it:

```
1[0-9]{2}|2[0-4][0-9]|25[0-5]
```

Tricky… tricky. The key is to find ways to match ranges, without missing any values or going outside the specified range. The common error was to miss values.

6. [16 points] Write an implementation of the C function that is specified below. Pay attention to the pre- and post-conditions, and the specified return value.

```
// Pre:
//          A[] contains AUsage values
// Post:
//          Barring an error, each element of A[] has been replaced
//          by its immediate successor, IF it divided its successor
//          evenly (i.e., without a remainder).
// Returns:
//          The number of elements that were replaced.
//
int Q6(int A[], int AUsage);
```

For example, given the arrays below, the function should respond as indicated:

```
A = {3, 9, 9, 2, 6, 12} → A = {9, 9, 9, 6, 12, 12} and returns 4
```

Here's one solution:

```
int Q6(int A[], int AUsage) {

   int numReplaced = 0;  // number of replacements that have been made

   // Trivial array traversal, EXCEPT that you must avoid looking for
   // a successor to the last element in the array:
   for (int idx = 0; idx < AUsage - 1; idx++) {

      // Subtlety:  you have to avoid dividing by zero:
      if ( A[idx] != 0 && A[idx + 1] % A[idx] == 0 ) {

         A[idx] = A[idx + 1];

         numReplaced++;
      }
   }

   return numReplaced;
}
```

DBZ: divide by zero   Lots of people (most) did not consider the possibility.

OOB: out-of-bounds   Most of those were off-by-one errors at the high or low ends.

From Discrete Math, x divides y iff dividing y by x yields a remainder of 0. The modulus operator, '%' just like in Java, yields the remainder.

7. Consider the following short C function. The given code, with appropriate include directives, compiles with no errors or warnings.

```
// Copies the contents of the array A[] into the arrays B[] and C[],
// alternating the elements.
//
// Example:
//                                                  B[]:  1  4  6  8  9 12
//    A[]: 1  3  4  5  6  7  8  9  9 11 12 -->
//                                                  C[]:  3  5  7  9 11
//
// Pre:      A[] contains ASZ elements in non-descending order.
//           B[] is large enough to hold ASZ/2 + 1 elements.
//           C[] is large enough to hold ASZ/2 + 1 elements.
// Post:
//           A[] is unchanged.
//           B[] and C[] contain alternating elements of A[],
//           as shown above.
//
void Q7(int32_t A[], uint32_t ASZ, int32_t B[], int32_t C[]) {

   uint32_t Apos = 0, Bpos = 0, Cpos = 0;

   while ( Apos < ASZ ) {          // walk A[] and copy...

      B[Bpos] = A[Apos];
      C[Cpos] = A[Apos + 1];

      Apos = Apos + 2;
      Bpos++;
      Cpos++;
   }
}
```

a) [6 points] The implementer of the C function shown above has assumed something that is not guaranteed by the preconditions. What is that?

The implementation copies pairs of elements from `A[]` into `B[]` and `C[]`.

That only makes sense if the number of elements in `A[]`, `ASZ`, is even.

There were lots of incorrect answers here, including (but not limited to):

- **A[] might be too large for its size to be specified by a uint32_t value; if so, you've got bigger problems, since an array that large will require 16GB of memory.**
- **One of the parameters passed to the function might be of the wrong type; e.g., the actual parameter corresponding to ASZ might be a signed int; if so, you've got the "caller is an idiot" error, but that will not automatically cause an error, since the caller's value will simply be mapped into the range of a uint32_t when it's passed.**
- **B[] or C[] might not be empty; nothing in the function specification says that any previous contents are supposed to be preserved, in fact the second post-condition and the example should make it clear that any previous contents will be over-written.**
- **The implementation assumes the values in A[] are in ascending order, not non-descending order; that's simply false, although the example is in ascending order.**

b) [8 points] Assuming all of the given preconditions are satisfied, how might this assumption affect the execution of the given function?

The preconditions say nothing about whether `ASZ` is even or odd; if `ASZ` is odd, then on the last pass through the loop, the code will attempt to access a memory location just past the end of the array.

The actual effect at runtime depends on ownership of that memory location; but, if there's not a runtime error, then the code will copy an unknown value into `C[]`.

The common issue here was that three things need to be said:

- **an array-out-of-bounds error will occur if ASZ is odd**
- **that could lead to a segmentation fault at runtime**
- **if not, it will lead to the copying of a value into C[] that should not go there**

And remember, there is no such thing as automatic bounds checking for arrays in C, nor is there any such thing as an "array index out of bounds" exception (or any other exceptions) in C.