

Exercises serial debugging

These exercises were originally developed for the Cyberinfrastructure Tutor, the web-based training site for High Performance Computing and Cyberinfrastructure topics:

<http://ci-tutor.ncsa.illinois.edu>

which is hosted by the National Center for Supercomputing Applications (NCSA). The exercises and solutions are adapted to show the use of different debugging tools and compiler options besides using traditional debuggers.

Please note that using a different compiler or debugger — for that matter even a different version of the same compiler or debugger — could result in a somewhat different behavior of the debugger.

Exercise 1: Checking variable values with the GDB debugger

The C code `variablePrinting.c`, is a very simple program that:

1. generates a 10-element array,
2. passes this array to a function that squares each element in the array and stores these values in a second 10-element array, and then
3. computes the difference between the element values in the two arrays.

We use this sample code to show an example debugging session using GDB to examine variable values.

After compiling and running the sample code we obtain the following output:

```
$ gcc -o variablePrinting variablePrinting.c
```

```
$ ./variablePrinting
```

```
array1 = [ 2 3 4 5 6 7 8 9 10 11 ]
```

```
The difference in the elements of array2 and array1 are: del = [ 0 0 0 0 0 0 0 0 0 0  
0 0 ]
```

The result for the array `del` certainly does not look correct.

a. What are the expected values for array `del`?

To identify this error we recompile our code using the debugging option `'-g'` and analyze it using GDB. In this example we illustrate how to debug a code by using a debugger to print out the values of selected variables during the execution of a program.

One possibility that could account for the values of all elements of `del` being zero is a coding error in our function `squareArray()`. If, for some reason, this function fails to square the elements of `array1` and, instead, is simply returning an array, `array2`, whose elements are identical to those of `array1` then

$$\text{del}[k] = \text{array2}[k] - \text{array1}[k] = \text{array1}[k] - \text{array1}[k] = 0$$

for all k .

Therefore, we might first want to examine the values of the elements of the array `array2` to see if $\{a_{2k}\} = \{a_{1k}\}$. We can do this in either of two ways. The first way is to set a breakpoint on the line of our code where we call the function `squareArray()` then print out the value of `array2[indx]` as we step through all `nelem` iterations of the `for()` loop within the body of this function.

```
int squareArray(const int nelem_in_array, int *array)
```

```

    {
        int indx;

        for (indx = 0; indx < nelem_in_array; indx++)
        {
            array[indx] *= array[indx];
        }
        return *array;
    }

```

The second way is to set a breakpoint immediately after the call to `squareArray()` and print out the elements of `array2` returned by the call to `squareArray()`. This example will illustrate both ways of examining these values $\{a_{2k}\}$.

After recompiling our code we run it using GDB.

```
$ gcc -g -o variablePrinting variablePrinting.c
```

```
$ gdb -tui ./variablePrinting
(gdb)
```

The `tui` option shows the source code in the upper half of the terminal, which is great to see where the execution of the program is, what will be the next command and to identify locations of breakpoints.

If we want to step into and through our function `squareArray()`, we need to set our initial breakpoint on line 37 and then rerun our code.

```
(gdb) b 37
Breakpoint 1 at 0x8048469: file variablePrinting.c, line 37.
(gdb) run
Starting program: variablePrinting
```

```
    array1 = [ 2 3 4 5 6 7 8 9 10 11 12 13 ]
```

```
Breakpoint 1, main () at variablePrinting.c:37
37      squareArray(nelem, array2);
```

We can then step into this function by issuing the `'step'` (or `'s'`) command.

```
(gdb) s
squareArray (nelem_in_array=12, array=0x80498d8) at variablePrinting.c:68
68      for (indx = 0; indx < nelem_in_array; indx++)
(gdb) s
70      array[indx] *= array[indx];
```

At this point we begin printing out the values of the elements of `array` on both the left- and right-hand side of the expression on line 70. We also track our progress thru the `for()` loop by printing out the value of the variable `indx` during each iteration of the loop. We can print the value of each of these variables by issuing the `'print'` (or `'p'`) command followed by the name of the variable whose value you want to print; e.g.,

```
(gdb) p indx
(gdb) p array[indx]
```

The `print` command

```
p v
```

prints out the value of the variable `v` in the same format as a `printf()` statement in C, that is

```
printf("%d\n", v)
```

if v is an integer variable or

```
printf("%f\n", v))
```

if v is a double.

In this example, we are only printing out the values of two variables. However, as the number of variables becomes more than just a few, this method of printing a sequence of variable values will begin to require a lot of typing. Fortunately, GDB provides alternatives for printing out a series of variable values without having to type multiple print commands each time you want to examine these values. One method involves issuing the GDB 'display' (or 'disp') command.

Let's display the loop variable and the array value

```
(gdb) disp indx
1: indx = 0
(gdb) disp array[indx]
2: array[indx] = 2
```

Now we can continue stepping through the loop.

It appears that our function `squareArray()` appears to be behaving properly.

Since each of the variables are automatically printed every time the program pauses, it is most useful only when examining variables whose values are changing frequently; e.g., during the execution of a `for()` loop. Otherwise, we would probably not want to see a set of constant values being printed out over and over throughout an entire debugging session. Fortunately, there is an easy way to tell GDB to stop displaying a given variable. We simply invoke the command 'undisp' followed by the number n that is printed to the left of the variable name,

```
n: variable-name = value
```

For example, we can cancel the printing of the variables `indx` and `array[indx]` by issuing the commands

```
(gdb) undisp 1
(gdb) undisp 2
(gdb)
```

Another option is to set a breakpoint in a longer loop and run, so you see the values only at that point for every iteration.

Examining the values of `array[indx]` we see that each element of `array2` returned by `squareArray()` is equal to the square of the corresponding element in `array1`. Consequently, we know that our coding error is not located anywhere within the function `squareArray()`. Therefore, we need to continue debugging our program until we locate the coding error in our program.

Complete the debugging of this code and identify the coding error in the program.

(BONUS e. How can you print the whole array `array1` in one go?)

Exercise 2: arrays

The default range of indices differs among the major programming languages. For example, in Fortran 77 the default index range is $i=1, \dots, N$ (where N is the array size); in C/C++ the default range is $i=0, \dots, N-1$; and in Fortran 90 the allowed index range can be *any* sequence of integers with *any* spacing. So if you routinely use several programming languages, it is easy to confuse the indexing rules.

This program is supposed to fill an integer array N elements according to an algorithm based on the mod operator (% in C, mod() in Fortran) and should add the values of all of the array elements with even indices and output the result. It should also do the same with odd-indexed elements. The program is written in two versions, both C and Fortran, just choose one.

In this exercise we use the Intel compiler. After compiling and executing the code, we get the following result:

```
salomon$ module load intel
salomon$ icc -o array array.c
salomon$ ./array
oddsun=5
evensun=224683
salomon$
```

As you can see, the code does not generate compiler or runtime errors. However, something is wrong. By looking at the body of the first loop it looks like the array elements should be relatively small integers. So the value of *oddsun* is probably alright. But why is the value of *evensun* so large? The value of *evensun* is computed in our summation loop, so we suspect that the error occurs there. Therefore, we will concentrate our debugging effort on that section of code.

There are five basic steps we will take to identify the error along with a final step of exiting the debugger:

1. Recompile using the '-g' option
2. Run the debugger (you can use gdb-ia, TotalView or DDT, whichever you prefer)
3. Identify the breakpoint
4. Execute the code up to the breakpoint
5. Step through the summation loop

The first step shouldn't be a problem anymore. If you will use GDB, use the command gdb-ia, it is built by Intel to match its compiler. Let's put a breakpoint at the beginning of the summation loop.

a. What's the command to do that? How do we run the program up to the breakpoint?

Before analyzing the summation loop, it is good idea to check out some of the array elements to see if they have the correct values.

b. Use the 'print' debugger command to show the values of the elements *tock*[2] and *tock*[7]. Are the values what you expected?

c. How do we display the value of *evensun* to make sure that it did get initialized to zero and how it gets updated?

Step through the summation loop until you see the problematic values.

d. What is the cause of the high (and random) *evensun* and how can you solve this?

e. How can you print the whole array *tock* in one go?

Exercise 3: learning valgrind

a. use valgrind to run the program variablePrinting from the first exercise. Valgrind finds two problems in our code, which ones? On which lines did these problems occur? Does this match the errors that you found in the first exercise?

c. use valgrind to run the program array from the second exercise. Does it detect the use of uninitialised values in your arrays? Although we only had one undefined value, valgrind seems to find many more. Why is that? What is the origin of the uninitialised values? See also <http://valgrind.org/docs/manual/mc-manual.html#mc-manual.uninitvals>

Exercise 4: Dynamic summation

The program `readsum.c` has an indexing error in it. The code reads in data from a file called `input.txt` and puts it in the array `tock`. Then the sum of all the elements is calculated and printed out.

The code compiles, runs, and even produces the following output:

```
$ gcc -o readsum readsum.c
$ ./readsum < input.txt
sum is 1090515646
$
```

As you can see from the output, the value of the sum is way too large given the numbers input. We leave it up to you to figure out why.

Hint: to run an application in `gdb` that reads the file `input.txt` from standard input, use

```
$ gdb ./readsum
(gdb) start < input.txt
```

What other tool could you use to analyze this application?

Exercise 5: Fortran I/O

For this lesson we use an example program `sections.f90` written in Fortran 90 that

1. reads in a set of data stored in ASCII format
2. does some simple data manipulation
3. writes the results as a binary file
4. reads the binary file back in before finally writing some of the fields to a formatted text file

As you will see, this program is very poorly written and full of errors.

When our sample code is compiled using the `ifort` compiler and run, it gives the following error:

```
$ ./sections
forrtl: severe (24): end-of-file during read, unit 10, file
/home/donnors/src/sections/sectionsCT.txt
```

There are two possible and common reasons why the `ifort` compiler gives this error message:

1. the program is trying to read more data than there is in the file
2. the file "sectionsCT.txt" does not even exist.

a. Which one is the problem here?

Fortran opens files for reading & writing by default and it creates the file if it doesn't exist. However, it's good practice to open files read-only when you don't intend to write to it. This prevents confusing error messages or accidentally overwritten files. So input-files should be opened only when they exist, and output files should be opened only if they don't already exist.

b. What is the optional argument to only open existing files? What is the optional argument to open the file read-only?

Edit the source code and run `'make clean; make'`. Now, when we run the program the runtime environment fails with the helpful message that the file does not exist.

Correct the file name in the source code, recompile and rerun the program. The Fortran runtime environment runs into more errors when reading the input file.

c. Compare the input that is read in the program to what you expect it should be reading (either through print statements or a debugger). When printing strings, it is often difficult to see if there are extra spaces at the beginning or end, how can you make these visible?

d. Edit the format to correctly parse the input-file.

Rerun the program and check that the output in the output file is as expected.

e. Is it correct? Please fix the rest of the program.

Exercise 6: argument and type checking

When a subroutine is called from within a program, the program call's argument list must match that of the subroutine it is calling. Since typically there are multiple items in a subroutine's argument list it can be easy to inadvertently create a mismatch between the subroutine's dummy argument list and the actual argument list passed by the corresponding program call. Some of the more frequent types of argument mismatches are:

- The number of terms in the argument list do not match
- The data types of the terms in the argument list do not match
- Data pass-by-value mismatch (When a C program calls a Fortran subprogram, all variables (including scalars) must be passed by reference instead of by value.)

Compilers react quite differently to these mismatches. Some may issue a warning while others may not. For C programs that use prototyping and Fortran 90 programs that declare subroutines via the f90 module interface, the compiler can readily perform an argument list compatibility check to see if there are any mismatches.

If you prefer Fortran, then rewrite the program `miss1.f90` to a Fortran-90 program using modules. Check your code with the Intel and the GNU Fortran compilers.

If you prefer C, then rewrite `miss1.c` to use prototyping.

Check that the compiler actually detects the incorrect use of arguments.

Exercise 7: C++ pointer arithmetic

Since the name of an array is simply an address (specifically, the address of the first element in the array), then we can always point to the beginning of an array of elements using a pointer. The statement

```
ptr = arr;
```

initializes a pointer, *ptr*, to point to `&arr[0]`, where *arr* is the name of some array.

If we want to write a function that returns the address of the beginning of the array *arr*, one way to do this is to have the function return the name of the array, *arr*. An alternative way would be to initialize a pointer, *ptr*, using the statement shown above and have the function return the pointer, *ptr*.

The C++ program `pointers.cc` uses this concept to copy the elements from one array to another array. Run this program to see if it completes successfully and, if not, debug the program in the debugger of your choice, identify the bug(s) in the code, and see if you can correct the error(s) so that the program runs successfully.

Note that in the code `pointers.cc` we are using the ANSI-C++ standard header files.

Compile and run the program. Choose your favorite set of tools to debug this problem:

- `addr2line`
- `valgrind`
- `gdb`
- `TotalView/DDT`

Exercise 8: real numbers

The Fortran code `trapezoid.f90` integrates the function $\cos(x)$ over the interval zero to pi using the trapezoidal rule. The result should be identically equal to zero. For demonstration purposes, only 10 intervals (11 endpoints) are used.

The sample code was compiled and run, and gave the following output:

i	x	f(x)	int(f(x))
1	0.00	1.00	0.00
2	0.00	1.00	0.00
3	0.00	1.00	0.00
4	0.00	1.00	0.00
5	0.00	1.00	0.00
6	0.00	1.00	0.00
7	0.00	1.00	0.00
8	0.00	1.00	0.00
9	0.00	1.00	0.00
10	0.00	1.00	0.00
11	3.14	-1.00	0.00

The four columns contain the endpoint number, the independent variable x , the dependent variable $\cos(x)$, and the integral of the function from the left endpoint to the current point. The good news is that the final value of the integral is exactly zero, as expected. The bad news is that most of the other values in the output file are clearly incorrect.

- Use your favourite debugger to find why the value of x does not advance through the loop.
- Could the GNU compiler be of help to locate these issues for you?

Exercise 9: dynamic memory

Static memory allocation has been the traditional method for assigning memory to an array for quite some time. In this method, memory is allocated at the beginning of the program and does not change for the duration of the program. One problem with this method is that the size of the array may not be known at the initial declaration. Sometimes the array size will be input from a file or user prompt or perhaps be calculated in the program. In these cases, memory cannot be statically allocated to the array.

Dynamic memory allocation is the ability to assign memory to an array at any point in a program. In dynamic memory allocation, the memory size can be determined from sources such as a file, user prompt, or calculation and memory of the appropriate amount is connected to the array name at the point in the code when it is needed. Unlike static memory allocation, this method does not tie up all the memory allocated to it throughout the duration of the program but rather just from the point at which it is allocated. In addition, the size declared is the exact size needed for the array and the memory can be released when the array is no longer needed.

Several high-level languages provide the ability to dynamically allocate memory. These languages typically provide this capability through built-in routines that can be called when memory needs to be allocated. The languages most commonly used for dynamic memory allocation are C, C++, and Fortran 90.

As with any programming technique, when dynamic memory allocation is performed incorrectly it results in a variety of errors. Also, how the error manifests itself depends on which compiler, loader, and machine is used. In one situation, the compiler could catch the programming error and in another an executable could be made. In the latter case, when the buggy code is executed either a run-time error occurs or incorrect results are produced.

Two of the most common programming errors in dynamic memory allocation are trying to

allocate too much memory and not allocating the memory at all. Also common are allocating a zero size, the incorrect size, and for multi-dimensional arrays not allocating enough memory for all the dimensions. This is by no means an exhaustive list of the errors that can occur but it should give you an idea of the most common ones.

Try the program `dynamic.f90` on an Intel computer using both the Intel and GNU compilers.

While the executable compiled with the GNU compiler gives a segmentation fault, the Intel compiler gives you the wrong result.

a) Can you find the problem in the executable compiled with the Intel compiler with a memory debugging tool like `valgrind`? Why do you think is that?

Of course, the incorrect result is quite obvious here, but you might not notice it directly in a large simulation.

b) What optional flag could you use with the Intel compiler to check for incorrect use of arrays?

c) With the knowledge in mind that different compilers could behave quite differently when compiling and running the application, how could you improve the robustness of your application?

Exercises parallel debugging

These exercises were originally developed for the Cyberinfrastructure Tutor, the web-based training site for High Performance Computing and Cyberinfrastructure topics:

<http://www.citutor.org>

which is hosted by the National Center for Supercomputing Applications (NCSA). The exercises and solutions are adapted to show the use of different debugging tools and compiler options besides using traditional debuggers.

Please note that using a different compiler or debugger — for that matter even a different version of the same compiler or debugger — could result in a somewhat different behaviors in the debugger.

Exercise 10: outsourced

Our sample program, `outsorce.c`, is written for two processes: by design, process 0 defines an array `sdata`, then sends it to process 1. Subsequently, process 1 receives the message from process 0, determines its size and then passes the information on to the function `do_work`, which searches and prints the maximum value.

Compile & run it. It doesn't find the expected value of 49, but instead it finds a value that looks suspect like an uninitialized value. Valgrind could be used to search for such an issue.

Although Valgrind is not primarily a parallel debugger, it can be used to debug parallel applications as well. When launching your parallel applications, prepend the `valgrind` command. For example:

```
$ mpirun -np 4 valgrind outsorce
```

The default version without MPI support will however report a large number of false errors in the MPI library, such as:

Salomon contains two Valgrind versions with MPI included to check for these errors. Load one of the following modules before compiling:

```
Valgrind/3.11.0-intel-2015b for Intel MPI  
Valgrind/3.11.0-foss-2015b for OpenMPI
```

Now run the program using valgrind:

```
export LD_PRELOAD=$EBROOTVALGRIND/lib/valgrind/libmpiwrap-amd64-linux.so  
mpirun -n 2 valgrind ./outsorce
```

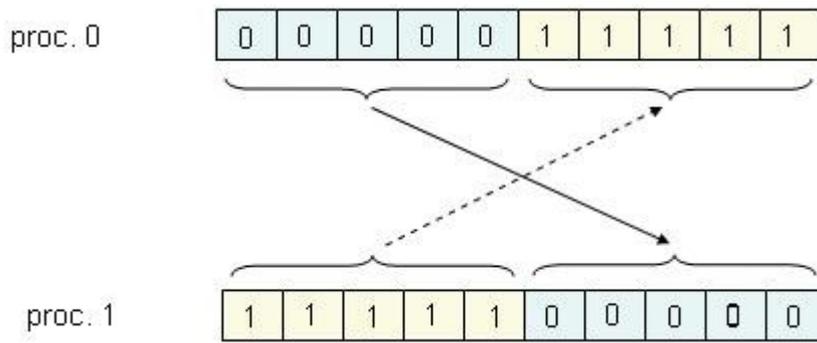
Which variables in our own program are not initialized?

b) Fix the program and check that it gives the expected result.

Exercise 11: crossover

The code `crossover.c` is designed to run on exactly two processors. An array is filled with process numbers. The first half of the array is filled with the local process number, and the second half of the array is filled with the other process number. The second halves of the local arrays are filled by message passing.

A diagram that depicts how the arrays are filled is shown below. Each processor fills the left half of the local copy of the array with its own process number. It then passes these values to the other processor, filling the right half of the local array.



To run the code, a command-line argument is required specifying the size of the one-dimensional array.

a) Compile and run the code using:

```
module load c openmpi
```

```
mpicc -o crossover crossover.c
```

```
mpirun -n 2 ./crossover
```

Try different sizes of the array up to at least 10,000 elements. How does the behaviour of the program change? Why do you think is that?

b) If you like, attach TotalView to the running job and try to see what's the problem.

c) Fix the program so that it also works for larger arrays.

Exercise 12: OpenMP

OpenMP can be used to parallelize code through the use of multiple threads on shared-memory systems. Shared memory means that a single address space is used for all threads, so threads can potentially write to memory locations simultaneously or in a non-deterministic order. In order to eliminate these occurrences, variables in parallel regions must be declared either *private* or *shared*. When the parallel region is entered, a separate copy of each private variable is created for each thread. The value of the variable can be changed independently on each thread, since each thread only operates on its own copy. The scope of the private copies of these variables is restricted to the parallel region unless directives are included to modify their dispensation upon entering (firstprivate) or leaving (lastprivate) the parallel region. A common OpenMP error is an incorrect declaration, resulting in the variable having the incorrect scope.

The Fortran code [simjet.f90](#) calculates the velocity profile of a two-dimensional planar jet of air, given two parameters defining the initial jet conditions. (Details of the physics are not of consequence here.) The output file, simjet.d, contains the parameters defining the jet and the velocity profile. At the end of the file is an edge momentum factor.

After compiling the code and running it serially, the end of the resulting output file was as follows:

```
1.105E+01  4.8600  4.364E-09  -1.373E-05  2.403E-04  -7.563E-01
1.153E+01  5.0672  2.883E-09  -1.374E-05  1.588E-04  -7.569E-01
1.202E+01  5.2827  1.872E-09  -1.375E-05  1.031E-04  -7.573E-01
edge momentum factor = 1.031E-04
```

An OpenMP directive was then added to parallelize the main loop as follows:

```
!$omp parallel do private(xi,txi,xmom)
```

```
do j = 1, nyd
```

```

xi = cxi*y(j)
txi = tanh(xi)
xmom = 1.0 - txi**2
u(j) = cu*xmom
v(j) = cv*(2.0*xi*xmom - txi)
enddo

```

This code was compiled and run with two threads resulting in the following end of the output file:

```

1.105E+01  4.8600  4.364E-09  -1.373E-05  2.403E-04  -7.563E-01
1.153E+01  5.0672  2.883E-09  -1.374E-05  1.588E-04  -7.569E-01
1.202E+01  5.2827  1.872E-09  -1.375E-05  1.031E-04  -7.573E-01
edge momentum factor =  0.000E+00

```

The velocity profile (table of numbers) is identical between the serial and parallel cases, but the edge momentum factor differs. Upon examining the entire output files, you would see that the whole file is identical between the two cases with the exception of the edge momentum factor.

Exercise 13: prefix

Introduction

In MPI programming, a common reason for a program to hang is a faulty assumption about the number or configuration of the MPI processes spawned by the parallel program. This is especially common when using MPI's *blocking* point-to-point communication routines (eg. MPI_SEND() and MPI_RECV()).

Objectives

In this lesson, you will learn how to diagnose programming errors involving incorrect process spawning assumptions and debug them using a parallel debugger.

The C program, [ring.c](#), constructs a ring topology out of its MPI processes and has each process send data to its neighbor on one side and receive data from its neighbor on the other side. However, as currently written, this program only works if it has four MPI processes. Find a way to make it work with an arbitrary number of processes. (Hint: there is more than one way to do this.)

Consider the binary tree example written in Fortran 77, prefix.f

This program constructs a binary tree topology among the MPI processes, which is used to pass around a series of tokens used to execute an arbitrary associative operation op(). The functions parent(), child0(), and child1() are used to determine the ranks of a process's neighbors in the tree structure. The program uses MPI_ISSEND() and MPI_RECV() to transport the operands for the associative operation, using different message tag values to indicate if a process should continue processing data or stop execution after the current operand.

The program also has two LOGICAL functions, is_leaf() and is_down(), used to determine a given MPI process's relative location in the tree and whether or not it has children with which to communicate.

As you would suspect, this program has a bug in it — when run, it hangs. For example, it displays the following behavior when running on 8 processors:

```

This is leaf          5 my starting value is          5
This is leaf          4 my starting value is          4
This is leaf          6 my starting value is          3

```

```
This is leaf          7 my starting value is          4
This is leaf          4 my final value is           4
This is leaf          5 my final value is           5
This is leaf          7 my final value is           3
This is leaf          6 my final value is           3
```

Note the lack of response from MPI processes 0 through 3. Worse, this behavior is independent of the number of MPI processes used.