

# Errata for ‘Using Advanced MPI’

May 9, 2017

- p 6** The formula in the computation for the number of neighbors in Conway’s Game of Life has an error in it. The second term on the third line should be  $u[i-1][j]$ . That is, the computation is

$$\begin{aligned} \text{nbrs} = & u[i+1][j+1] + u[i+1][j] + u[i+1][j-1] + \\ & u[i][j+1] + u[i][j-1] + \\ & u[i-1][j+1] + u[i-1][j] + u[i-1][j-1]; \end{aligned}$$

Thanks to Yang Shangqin.

- p 35** The size of the array `destinations` is  $\sum_i \text{degree}_i$ , not  $\sum_i \text{sources}_i$ .  
Thanks to Victor Eijkhout.

**Section 8.2 (Pages 243ff.)** This is not exactly an errata, but the section describes several of the functions for working with `MPI_Count`, but omits one of the most important, `MPI_Get_elements_x`, the `MPI_Count` counterpart of `MPI_Get_elements`. This routine is described on page 113 in the MPI 3.1 standard.

- p 256–264** This is not exactly an errata, but the book was published before the `MPI_T_pvar_get_index` and `MPI_T_cvar_get_index` routines were officially adopted by the MPI Forum into MPI. Until those routines become part of MPI, it is necessary to search through all defined control or performance variables by index to find the index that matches a name. Code for that is available on the web site, in the examples for the chapter “Support for Performance and Correctness Debugging.”

- p 130–134** Some readers have observed problems in using the MCS locks, described on pages 130–134. The code as presented is correct, but correct operation requires that the MPI implementation provide what is called asynchronous progress. This is required by the MPI standard; however, providing this may have negative performance consequences. Because of this, many MPI implementations by default do not enable asynchronous progress; some implementations may not provide a full implementation of asynchronous progress.

To use the MCS lock code, make sure that asynchronous progress is enabled in your MPI implementation. For MPICH, for example, this is done

by setting the variable `MPIR_CVAR_ASYNC_PROGRESS=1` (and also building MPICH to support asynchronous progress). There is some evidence that some implementations still have problems with asynchronous progress, even when this is requested.

For implementations that do not correctly support asynchronous progress, the most common workaround is for each process to periodically call some MPI communication routine, such as `MPI_Iprobe`. In the case of the MCS lock release code, there is this block of code:

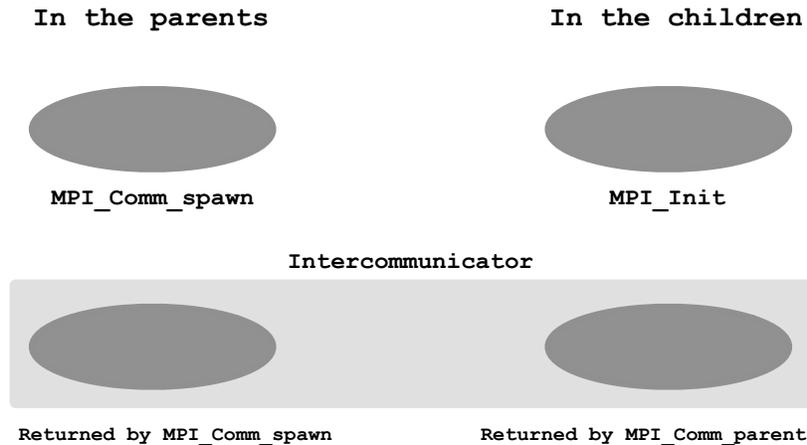
```

if (lmem[nextRank] == -1) {
    If-Block;
    // starts with MPI_Compare_and_swap(...)
}

```

The test in this code is used to avoid an unnecessary MPI communication operation; specifically, the `MPI_Compare_and_swap` call. However, because of the way that this code is written, the operation of the code is correct if the “If-Block” is unconditionally executed. In addition, the compare-and-swap call serves as an MPI communication call, so if asynchronous progress is not supported, unconditionally executing the “If-Block” can help (though not guarantee, depending on timing) the code to work properly.

**p 273** Figure 10.2, on the right, has `MPI_Int` where it should have `MPI_Init` (`MPI_Init_thread` would also be correct there). The figure should be:



and is available for download as [advmpi-spawn.pdf](#).

Thanks to Jeff Hammond.