

Suggested Solutions – Exercises of Lecture 9 – MVP'06

Alexandre David

Exercise 6.2

The restructuring of the application has been discussed during the lecture. You will need a buffer if you want to continue to work on your data, which means add a copy before sending. The sender has the advantage that it will not block and the disadvantage that it will have the overhead of a copy (still cheaper than waiting on the communication operation to terminate). The receiver has to block if the data has not arrived. So all depends on the amount of computation between communication operations and how the sender and receiver are synchronized.

Exercise 6.3

MPI allows for the two kinds of blocking and buffered *MPI_Send* implementations since they have the same semantics from a programmer point of view. So there is no reason to forbid having these two. Allowing for different implementations opens the possibilities of optimizing the communication depending on the kind of messages we have to send, in particular it is cheap to buffer small messages but more expensive (in time and memory) for large messages. In addition, depending on the underlying architecture, it may be better to never use buffering, e.g., if running on a SMP machine with shared memory.