Abstract

There are many works that consider the use of a spanning tree to perform collective operations in distributed environments. In most of them, the construction of that topology occurs at application startup, and posterior changes in the execution environment are not considered. However, distributed systems are dynamic, presenting performance variations, failures and recoveries of communication channels. This work presents a tool that provides collective operations for MPI that consider such dynamic behaviour. At first, the tool builds a minimum spanning tree representing the initial topology and then performs adaptations according to the data collected by the Network Weather Service (NWS). The tree construction and adaptation algorithms are executed in a distributed fashion.

Keywords: Minimum spanning tree, MPI collective communication, NWS monitoring tool.