

Homework 4

Donghoon Kim
dkim2@ncsu.edu

An MPI failure detector over PMPI

1. Problem description

Modern scientific applications on large scale MPPs have execution times ranging from day to months. These long-running MPI applications on clusters are prone to node or network failures as the systems scale[1]. The MPI application may have no progress in the case of node or network failures if such an application needs to exchange its computation results through the communication. Furthermore, the recovery overhead would be increased unless the fault detection services provide timely detection. On the other hand, the overhead of fault detection would be increased as the frequency of fault detection increases for monitoring accurate failures. It is still necessary to implement synthetic transparent fault detector on MPI standard for contemporary applications while there are many theoretical solutions for fault detection services because of the property of unreliable failure detectors, that is, completeness and accuracy[2].

2. Outline for approaching a solution

Assuming that the system model provides certain temporal guarantees on communication or computation called partially synchronous [3], the Fault Detector (FD) is able to utilize time based scheme, namely, ping-ack based implementation. The FD is a thread created by `pthread_create()` function which works independently under the application program. The FD uses two messages, ALIVE and ACK. ALIVE message is to check whether a destination node is alive or not. ACK message is to verify from a destination node. The FD should consider the time delay between two nodes with communication and computation time. The FD could suspect a destination node to be failed if no ACK message is received in correspondence to ALIVE message. The FD should be integrated into MPI environment with the following two properties:

- Transparency – The FD is launched in `MPI_Init()` routine with a profiling interface, which creates FD threads. The FD runs independently with a unique communicator different from an application program.
- Scalability – The FD sends a check message sporadically at any time when an application program has a routine to communicate. It would not lead to high communication overhead compared with the frequency of periodic check message since the FD at each node avoids redundant check messages for a defined time period.

References

- [1] Jitsumoto, H., Endo, T., Matsuoka, S., "ABARIS: An Adaptable Fault Detection/Recovery Component Framework for MPIs," *IEEE International Parallel and Distributed Processing Symposium (IPDPS 2007)* pp.1-8, March 2007.
- [2] Tushar Deepak Chandra , Sam Toueg, Unreliable failure detectors for reliable distributed systems, *Journal of the ACM (JACM)*, v.43 n.2, p.225-267, March 1996
- [3] Srikanth Sastry, Scott M. Pike , Jennifer L. Welch "Crash fault detection in celerating environments" *IEEE International Parallel and Distributed Processing Symposium (IPDPS 2009)* pp.1-12, 2009.