**Abstract**

Common cluster computing applications (like distributed databases and PVM applications) use the standard TCP/IP as their communication protocol. For a relatively error-free cluster environment, such general purpose protocols lead to an underutilization of the network bandwidth and consequently, a degradation in performance.

The bottleneck for faster communication are the protocol overheads namely

- *Data touching operations :* data copying and computing checksums.
- *Per packet interrupt overhead.*
- *VM (Virtual Memory) overheads :* due to allocation and deallocation of buffers.
- *Software latency :* due to the layers of protocol code the data has to undergo in the tcp protocol stack.

The primary objective of the QMP (Quick Message Protocol) is to provide network support tailor-made for the cluster environment. Each of the above overhead issues was adressed in some form or the other in the QMP protocol. We evaluated QMP's performance by running distributed benchmarks using TCP and then using the QMP protocol suite. The QMP coexists with the Tcp/Ip protocol stack. All results were obtained on 233 MHz Intel Pentium IIs running Red Hat Linux 6.2 with Compex 10/100 Mbps compatible network adapters.