

Extending τ -Lop to model concurrent MPI communications in multicore clusters

Achieving optimal performance of MPI applications on current multi-core architectures, composed of multiple shared communication channels and deep memory hierarchies, is not trivial. Formal analysis using parallel performance models allows one to depict the underlying behavior of the algorithms and their communication complexities, with the aims of estimating their cost and improving their performance.

LogGP model was initially conceived to predict the cost of algorithms in mono-processor clusters based on point-to-point transmissions with network latency and bandwidth based parameters. It remains as the representative model, with multiple extensions for handling high performance networks, covering particular contention cases, channels hierarchies or protocol costs. These very specific branches lead LogGP to partially lose its initial abstract modeling purpose.

More recent lognP represents a point-to-point transmission as a sequence of implicit transfers or data movements. Nevertheless, similar to LogGP, it models an algorithm in a parallel architecture as a sequence of message transmissions, an approach inefficient to model algorithms more advanced than simple tree-based one, as we will show in this work.

In this paper, $\tau\tau$ -Lop model is extended to multi-core clusters and compared to previous models. It demonstrates the ability to predict the cost of advanced algorithms and mechanisms used by mainstream MPI implementations, such as MPICH or Open MPI, with high accuracy. $\tau\tau$ -Lop is based on the concept of concurrent transfers, and applies it to meaningfully represent the behavior of parallel algorithms in complex platforms with hierarchical shared communication channels, taking into account the effects of contention and deployment of processes on the processors. In addition, an exhaustive and reproducible methodology for measuring the parameters of the model is described.

Fuente de la publicación:

- Juan-Antonio Rico-Gallego, Juan-Carlos Díaz-Martín, Alexey L. Lastovetsky.
[Extending \$\tau\$ -Lop to model concurrent MPI communications in multicore clusters](#) [1].
Future Generation Computer Systems, Volume 61, August 2016, Pages 66-82, ISSN 0167-739X,
<http://dx.doi.org/10.1016/j.future.2016.02.021> [2].

Proyectos relacionado:

- [Evaluación de AzequiaMPI](#) [3].

URL del envío: <http://www.cenits.es/enlaces/publicaciones/extending-t-lop-model-concurrent-mpi-communications-multicore-clusters>

Enlaces

[1] <http://www.sciencedirect.com/science/article/pii/S0167819115000447>

[2] <http://dx.doi.org/10.1016/j.future.2016.02.021>

[3] <http://www.cenits.es/proyectos/evaluacion-de-azequiampi>